



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s): Richard HULL, et al.) Group: 2681
Serial No.: 10/635,940) Examiner: Not yet assigned
Filed: August 5, 2003) Our Ref: B-5192 621141-3
For: "METHOD AND APPARATUS FOR)
PROVIDING INFORMATION ABOUT)
A REAL-WORLD SPACE") Date: January 5, 2004

CLAIM TO PRIORITY UNDER 35 U.S.C. 119

MAIL STOP MISSING PARTS
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

[X] Applicants hereby make a right of priority claim under 35
U.S.C. 119 for the benefit of the filing date(s) of the
following corresponding foreign application(s):

<u>COUNTRY</u>	<u>FILING DATE</u>	<u>SERIAL NUMBER</u>
GREAT BRITAIN	6 August 2002	0218188.1

[] A certified copy of each of the above-noted patent
applications was filed with the Parent Application
No. _____.

[X] To support applicant's claim, a certified copy of the above-
identified foreign patent application is enclosed herewith.

[] The priority document will be forwarded to the Patent Office
when required or prior to issuance.

I hereby certify that this correspondence
is being deposited with the United States
Postal Service with sufficient postage as
first-class mail in an envelope addressed
to the "Mail Stop Missing Parts, Commissioner
for Patents, P.O. Box 1450, Alexandria, VA
22313-1450", on January 5, 2004
by Suzanne Johnston

Respectfully submitted,

Richard P. Berg
Attorney for Applicant
Reg. No. 28,145

LADAS & PARRY
5670 Wilshire Boulevard
Suite 2100
Los Angeles, CA 90036
Telephone: (323) 934-2300
Telefax: (323) 934-0202



INVESTOR IN PEOPLE


The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated 28 July 2003

17



18

19

20

21

22



Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference

300204857-01 GB

2. Patent application number

(The Patent Office will fill in this part)

0218188.1

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Hewlett-Packard Company
3000 Hanover Street
Palo Alto
CA 94304, USA

Patents ADP number (*if you know it*)

Delaware, USA

If the applicant is a corporate body, give the country/state of its incorporation

49 6588 00 4

4. Title of the invention

METHODS AND ARRANGEMENTS APPLICABLE TO EXHIBITION SPACES

5. Name of your agent (*if you have one*)

"Address for service" in the United Kingdom to which all correspondence should be sent (*including the postcode*)

Robert Squibbs
Hewlett-Packard Ltd, IP Section
Filton Road, Stoke Gifford
Bristol BS34 8QZ

4077 442 002

Patents ADP number (*if you know it*)

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (*if you know it*) the or each application number

Country

Priority application number
(*if you know it*)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (*Answer 'Yes' if:*

Yes

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description 65

Claim(s) 17

Abstract 1

Drawing(s) 12

102

[Signature]

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

1

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

Fee Sheet

11. I/We request the grant of a patent on the basis of this application.

Signature

T. F. Squith

5/8/2002

Date

12. Name and daytime telephone number of person to contact in the United Kingdom

Tony Judd

Tel: 0117-312-8026

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

Methods and Arrangements Applicable to Exhibition Spaces

Field of the Invention

- 5 The present invention relates to methods and arrangements applicable to exhibition spaces; however, the methods and arrangements also have application in other contexts.

Background of the Invention

- 10 Ant colony optimization is concerned with the development and use of optimization algorithms inspired by the collective behaviour of large colonies of social insects. Typically, the algorithms represent a problem space as a network of nodes connected by arcs. The network is traversed by simple, autonomous agents that are capable both of depositing virtual markers on nodes and arcs, and acting upon the accumulated markers that they encounter on their travels. By an appropriate choice of agent behaviours, some
15 emergent characteristic of the network can be caused to converge on a (near-)optimal solution to the original problem. This approach has been successfully applied to a range of problems such as the 'Traveling Salesman' Problem and job scheduling. More detailed information about the approach can be found in *Swarm Intelligence* by Bonabeau, Dorigo & Teraulaz (Oxford University Press, 1999).

- 20 Artificial life seeks to understand the processes by which biological and social complexity arise from simple organisms or agents. Typically, the approach is to construct a computer simulation of a universe populated by such agents and to study their interaction and evolution. The simulated universes may or may not reflect natural laws, and the agents may
25 or may not be modeled on naturally occurring organisms. Within that context, the simulation of ant-like agents with the capability to deposit and sense virtual markers (pheromones) has been known for at least a decade (for example, see *Ant Farm: Towards Simulated Evolution* by Collins & Jefferson (in *Artificial Life II*, Farmer et al, Addison Wesley, 1991).

- 30 Agent-based robotics applies similar ideas to motivate the development and exploration of swarms of simple interacting robots operating in the real world. The idea of pheromone

deposition and detection is well known in this field but is primarily used metaphorically to inspire mechanisms that actually implement *direct* communication between individuals rather than *indirect* communication through the environment in which the individuals move. For example, see *Progress in Pheromone Robotics* by Payton, Estkowski & Howard
 5 (preprint, *7th International Conference on Intelligent Autonomous Systems*, March 25-27, 2002). An exception is the work of Andrew Russell in which robots do deposit (and sense) a chemical marker directly into the environment (see <http://www.ecse.monash.edu.au/staff/rar/>).

- 10 It is an object of a first aspect of the present invention to use the concept of pheromones to provide trail information about users of an exhibition space

In many mobile computing applications, there may be a requirement that users follow a particular path through a physical space. However, the physical space may be devoid of
 15 physical signs to indicate a specified path through that space. There are many uses of audio to guide navigation, including the use of audio beacons to attract users to its source, and the use of sonar to indicate obstacles ahead. A system of audio cues known as the "Oboe" system was also used in the Second World War to guide the pilots of RAF (the British Royal Air Force) bombers to targets; in this system monaural audio cues were presented to
 20 the pilot through headphones and represented three ternary states, namely: turn left, turn right, and straight ahead.

It is an object of second and third aspects of the present invention to provide sound based cues for guiding a user along a target path.

25

In many uses of mobile computing devices, data (and/or services) stored on networked servers are associated with particular physical locations. The mobile devices are expected to access that data via a network connection when they are at the locations associated with the data. There is a latency associated with accessing data over a network, owing to delays
 30 in the network and the server hosting the data.

It is an object of a fourth aspect of the present invention to minimize this latency.

Once data has been loaded into the memory of a mobile device, it is well known to retain the data beyond its initial use in order to speed subsequent accesses to the same data, such
 5 subsequent accesses being a common occurrence. As the device memory is of finite size, data held in the device must occasionally be removed, for example to make space for new data.

It is a further object of the fourth aspect of the present invention to provide a way of
 10 determining what data to flush from the memory of the mobile device.

It is an object of a fifth aspect of the present invention to free up memory space in a mobile device without flushing data items.

15 Device memory is but one limited resource in a system in which data items are provided from a service system to multiple mobile devices. The service system and/or the communications infrastructure between this system and the mobile devices can become a bottleneck if the mobile devices are all frequently requesting data items. It is known to provide distributed or cooperative caching in which data can be stored in and retrieved
 20 from any of multiple data caches. There are many algorithms for finding data in such caches and for maintaining cache coherency. For example, see *A Survey of Cooperative Caching*, Raunak, M. at <http://citeseer.nj.nec.com/432816.html>.

It is object of a sixth aspect of the present invention to provide a way of reducing the
 25 loading on the service system used to store data items of interest to mobile device users.

For members of a party visiting a space with which media objects are associated, the shared experience can be enhanced by the sharing of such objects by the group members. Conventionally, media objects are shared by reference, for example by passing an
 30 appropriate URI. Where the media object is a streamed object, a recipient using a shared reference to the media object will typically experience the streamed media object from its beginning, whilst the person who passed on the reference will already be some way

through the streamed object. However, the person who passed on the reference may wish the recipient of the reference to experience the media object in a synchronized manner, i.e. to ensure that they both experience the same parts of the media stream in the same order and at the same time. Colloquially, one person wishes to invite the second person to “listen to this” (or “look at this “ etc). Multicast streaming protocols are known which enable a single media stream to be sent to multiple devices over the Internet with synchronization of multiple media channels within a composite, structured media stream (e.g. SMIL). However such protocols are not widely deployed in the internet.

10 It is an object of a seventh aspect of the present invention to provide a way of coordinated presentation of a streamed media object on multiple devices without the use of multicasting protocols.

15 **Summary of the Invention**

According to one aspect of the present invention, there is provided a method of providing information about a real-world space, comprising the steps of :

- (a) as each of at least one first user moves through said space, a succession of virtual markers is deposited by a mobile device carried by the user in a digital representation of the space where they are stored to indicate locations visited by the user in the space;
- (b) on a second user moving through the space, the virtual markers deposited by the said at least one first user are accessed to provide to a mobile device of the second user information for facilitating use of the space by the second user.

25 In a preferred embodiment, the virtual markers each have an initial strength value when deposited and the markers are aggregated, in dependence on their associated locations, by aggregating their strength values by location either when stored or when used. Advantageously, the strength of the marker aggregations are caused to decay with time.

30 The markers associated with the first users can be used to provide path information about the most/least popular routes as well as information about the most popular items of interest in the space. The markers can also be used to help predict where the second user

might go next whereby to enable pre-emptive caching of media items that the second user may need in the near future.

The first aspect of the present invention also envisages an arrangement for implementing
5 the foregoing method.

According to a second aspect of the present invention, there is provided a method of guiding a user along a target path, comprising the steps of:

- (a) determining the position of the user relative to the target path; and
- 10 (b) providing respective audio cues to the user via left and right audio channels, these cues being indicative of the relative position determined in step (a) and varying in a complementary manner over at least a range of values of said relative position.

The second aspect of the present invention also envisages an arrangement for implementing the foregoing method.

15

According to a third aspect of the present invention, there is provided a method of guiding a user along a target path, comprising the steps of:

- (a) determining the position of the user relative to the target path; and
- (b) rendering an audio beacon through audio output devices carried by the user such that
20 the beacon appears to the user to emanate from a location in a direction at least approximating the direction of the target path onward from the user's current position.

The third aspect of the present invention also envisages an arrangement for implementing the foregoing method.

25 According to a fourth aspect of the present invention, there is provided a method of managing a cache of a mobile device carried by a user, the cache being used for storing items associated with locations in a real-world space being visited by the user; the method comprising the steps of:

- (a) determining the probability of usage of an item in dependence on the user's progress
30 around the space;
- (b) changing the contents of the cache by adding or removing an item on the basis of the determination carried out in step (a) in respect of that item or other items

The fourth aspect of the present invention also envisages an arrangement for implementing the foregoing method.

According to a fifth aspect of the present invention, there is provided a method of
 5 managing a cache of a mobile device carried by a user, the cache being used for storing items associated with locations in a real-world space being visited by the user; the method comprising the steps of:

- (a) receiving an item at the mobile device, and
- (b) degrading the received item to reduce the amount of cache space needed to store it, and
 10 storing the degraded item in the cache instead of the un-degraded item.

The fifth aspect of the present invention also envisages an arrangement for implementing the foregoing method.

According to a sixth aspect of the present invention, there is provided a method of
 15 retrieving a data item to a mobile device carried by a first user visiting a real-world space, the data item being available from a service system to mobile devices of users visiting the space; the method comprising the steps of:

- (a) seeking to retrieve the data item to the first user's mobile device by transfer from another mobile device in said space;
- 20 (b) in the event that step (a) is unsuccessful, retrieving the data item to the first user's mobile device by transfer from the service system.

The sixth aspect of the present invention also envisages an arrangement for implementing the foregoing method.

25 According to a seventh aspect of the present invention, there is provided a method of coordinated consumption of a streamed media object by first and second devices, the media object being accessible for streaming from a server, the method comprising the steps of:

- (a) streaming the media object from the server to the first device and presenting it to a user of this device;
- 30 (b) sending from the first device to the second device, during the course of step (a), data identifying the media object and a current position reached in presenting the object to the user of the first device;

(c) in response to a request from the second device, streaming the media object from the server to the second device in a separate stream to that involving the first device, and presenting the media object to the user of the second device such that normal presentation commences at a position at or with an advance relative to the said current position indicated in step (b).

The seventh aspect of the present invention also envisages an arrangement for implementing the foregoing method.

10 **Brief Description of the Drawings**

Embodiments of the invention will now be described, by way of non-limiting example, with reference to the accompanying diagrammatic drawings, in which:

- . Figure 1 is a diagram of an exhibition hall having an arrangement for delivering relevant media objects to visitors in a timely manner as the visitors encounter items of interest in the hall;
- . Figure 2 is a diagram of a mobile device and service system used in the Figure 1 arrangement;
- . Figure 3 is a diagram of a location report sent from the mobile device to the service system of Figure 2;
- . Figure 4 is a diagram of a response message sent by the service system to the mobile device of Figure 2;
- . Figure 5 is a diagram illustrating some of the choices available when implementing a pheromone trail mechanism such as provided by the mobile device and service system of Figure 2;
- . Figure 6 is a diagram depicting the functional blocks involved in providing a pheromone trail mechanism;
- . Figure 7 is a diagram showing a target path to be followed by the user using audio guidance sounds generated by a first embodiment of a path guide unit of the Figure 2 mobile device;
- . Figure 8 is a diagram showing variation in frequency with distance from the target-path centreline, of the audio guidance sounds produced by the first embodiment of the path guide unit;

- . Figure 9 is a state diagram of a specific implementation of the first embodiment of the path guide unit;
- . Figure 10 is a diagram showing how the control regime employed by the Figure 10 implementation of the first embodiment of the path guide unit, varies with
5 the angle of moving / facing of the user relative to the target-path centreline;
- . Figure 11 is a diagram showing a target path to be followed by the user using audio guidance sounds generated by a second embodiment of the path guide unit of the Figure 2 mobile device;
- 10 . Figure12 is a diagram showing, for a variant of the second embodiment of the path guide unit, the sounds produced by three virtual sound beacons to provide audio guidance to the user;
- . Figure13 is a diagram depicting a wake zone behind a user progressing around the Figure 1 hall;
- 15 . Figure14 is a diagram illustrating the fall-off in allocated item usage probability with distance along and to the side of a movement track;
- . Figure15 is a table showing for each of multiple virtual features located around the Figure 1 exhibition hall, the likely next feature to be visited from the current feature based on counts of what previous users have done;
- 20 . Figure16 is a flow chart illustrating the operation of a first implementation of an arrangement by which the Figure 2 mobile device seeks to obtain a desired feature item first from another mobile device before requesting the item from the service system;
- . Figure17 is a flow chart illustrating the operation of a second implementation of an
25 arrangement by which the Figure 2 mobile device seeks to obtain a desired feature item first from another mobile device before requesting the item from the service system; and
- . Figure18 is a diagram illustrating coordinated consumption of a streaming media object by two mobile devices.

Best Mode of Carrying Out the Invention

Figure 1 depicts a real-world environment for which a number of zones have been defined in a virtual world that maps onto the environment. When a person moving in the environment (called a “user” below) is detected as moving into one of these zones, one or more media objects are delivered to the user via a communications infrastructure and a mobile device carried by the user. A zone may correspond to an area around a real-world object of interest with the media object(s) delivered to a user in this area relating to that real-world object. Alternatively, a zone may not correspond to any real-world object .

- 10 In considering such an arrangement, it is convenient, though not essential, to introduce the abstraction of a virtual feature which is the subject of each zone. Each such virtual feature is given a number of properties such as a unique identifier, a location in the real-world environment, the real-world extent of the zone associated with the feature, a subject description indicating what the feature concerns, and a set of one or more media-object
- 15 identifiers identifying the media objects (or “feature items”) associated with the feature. The zone associated with a virtual feature is referred to hereinafter as the ‘active zone’ of the feature.

- For a feature that is intended to correspond to a particular real-world item (and typically having an active zone that maps to an area about a real-world object), this can be indicated in the subject description of the feature. Using the feature abstraction makes it easier to associate feature items that all relate to the same zone and also facilitates adding / removing these features items since data about the real-world extent of the related zone is kept with the feature and not each feature item.

25

Each feature is represented by a feature record held in a data-handling system, the feature records together defining the aforesaid virtual world that maps to the real-world environment. Each feature can be thought of as existing in this virtual world with some of these virtual features mapping to real-world objects.

30

As already noted, when a user is detected as within an active zone of a feature, one or more feature items are delivered to the mobile device of the user for presentation to the user. A

feature item can be presented automatically to the user upon delivery or the item can be cached and only presented upon the user having expressed an interest in the feature in some way such as by dwelling in the active zone of the feature more than a minimum time or by explicitly requesting presentation of the feature item. Indeed, the delivery of the feature item to the mobile device can also be deferred until the user is detected as having expressed an interest in the feature; however, since this approach can introduce a delay before the item is available for presentation, the embodiments described below deliver feature items to the mobile device of the user without awaiting a specific expression of interest in each feature (though, of course, a general filtering may be applied as to what items are delivered according what types of features are of interest to the user). Preferably, each feature or feature item is given a property indicating whether feature item delivery is to be effected automatically upon delivery or only after a user has expressed an interest in the feature; this enables important items (such as warning messages concerning features associated with potentially hazardous real-world items) to be pushed to the user whilst other items are subject to an expression of interest by the user. Advantageously, a user may elect to have feature items automatically presented even when the corresponding feature/item property does not require this. Furthermore, since as will be described hereinafter, pre-emptive caching of feature items in the user's mobile device may be implemented, automatic presentation is qualified so as only to apply where the user is in the active zone of the feature with which the feature item is associated.

Considering the Figure 1 example in more detail, the environment depicted is an exhibition hall 10 having rooms 11 to 17 where:

- room 11 is an entrance foyer with reception desk 18 but no associated virtual features;
- room 12 is a reference library with no associated virtual features;
- rooms 13, 14 and 15 are used for displaying real-world objects, namely paintings and sculptures, for each of which there is a corresponding virtual feature centred on the object concerned and with an associated active zone (indicated by a dashed line);
- room 16 is used for experiencing virtual features for which there are no corresponding real-world objects, the location associated with each feature being

- indicated by a cross 22 and the corresponding active zone 25 by a dashed line; and
- room 17 is a cafeteria with no associated virtual features.

Virtual features are also defined in correspondence to the majority of openings 23 between rooms, the active zones 25 associated with the features again been indicated by dashed lines. Typically, the feature items associated with these features are incidental information concerning the room about to be entered and are automatically presented. It will be seen from Figure 1 that only a single feature is applied to an opening 23 so that it is not possible to tell simply from the fact that a user is detected in the active zone of the feature which room the user is about to enter; however, as will be later described, it is possible to determine from the user's past activity (either location based or feature based) the general direction of progression of the user and therefore which room is about to be entered. This enables the appropriate feature item to be selected for delivery to the user from amongst the items associated with the feature.

On entering the exhibition hall 10, a user 30 collects a mobile device 31 from the reception desk 18 (or the user may have their own device). This device 31 cooperates with location-related infrastructure to permit the location of the user in the hall 10 to be determined. A number of techniques exist for enabling the location of the user to be determined with reasonable accuracy and any such technique can be used; in the present example, the technique used is based on an array of ultrasonic emitters 33 (represented in Figure 1 by black triangles) positioned at known locations in each room (typically suspended above human level). The emitters 33 are controlled by controller 32 to send out emitter-specific emissions at timing reference points that are indicated to the mobile device 31 by a corresponding radio signal sent by the controller 32. The device 31 is capable of receiving both the timing reference signals and the emissions from the ultrasonic transmitters 33. The device 31 is also pre-programmed with the locations of these emitters and is therefore able to calculate its current location on the basis of the time of receipt of the emissions from the different emitters relative to the timing reference points.

The exhibition hall is equipped with a wireless LAN infrastructure 36 comprising a distribution system and access points 37. The wireless LAN has a coverage encompassing substantially all of the hall 10, the boundary of the coverage being indicated by chain-

dashed line 38 in Figure 1. The wireless LAN enables the mobile device to communicate with a service system 35 to download feature items appropriate to the feature (if any) corresponding to the current location of the user. In the present example, the determination of when the location of the user (as determined by the device in the manner already described) places the user within the active zone of a virtual feature, is effected by the service system; however, it is also possible to have the device 31 carry out this determination provided it is supplied with the appropriate information about the feature zones.

- 10 It will be appreciated that communication between the device 31 and service system 35 can be effected by any suitable means and is not limited to being a wireless LAN.

Figure 2 shows the mobile device 31 and service system 35 in more detail. More particularly, the mobile device 31 comprises the following functional blocks:

- 15 - A location determination subsystem 40 with an associated timing reference receiver 41 and ultrasonic receiver 42 for receiving the timing reference signals from the location infrastructure 32 and the emissions from the ultrasonic emitters 33 respectively; the location determination subsystem 40 is operative to use the outputs of the receivers 41 and 42 to determine the location of the mobile device (as already described above) and to send location reports to the service system 35.
- 20 - A visit data memory 43 for holding data about the current “visit” – that is, the current tour of the hall 10 being undertaken by the user of the mobile device 31.
- A feature-item cache 44 for caching feature items delivered to the mobile device 31 from the service system 35. The cache 44 has an associated cache manager 45.
- 25 - A communications interface 46 for enabling communication between the mobile device 31 and the service system 35 via the wireless LAN infrastructure 36.
- A user interface 48 which may be visual and/or sound based; in one preferred embodiment the output to the user is via stereo headphones 60.
- A visit manager 47 typically in the form of a software application for providing control and coordination of the other functions of the mobile device 31 in accordance with input from the user and the service system 35.
- 30 - A visit path guide 49 for giving the user instructions / indicators for following a

planned route around the hall 10.

Much of the foregoing functionality will typically be provided by a program-controlled general purpose processor though other implementations are, of course, possible.

- 5 The visit data held by memory 44 will typically include a user/device profile data (for example, indicating the subjects of interest to the user, the intended visit duration, and the media types that can be handled by the device), an electronic map of the hall 10, the user's current location as determined by the subsystem 40, and the identity of the feature (if any) currently being visited together with the IDs of its related feature items. The visit data also
- 10 includes a feature history for the visit, which is either:
- the history of visited features and their related feature item IDs in the order the features were visited (thus, a feature is added to the top of the visited-feature history list when the feature is encountered), or
 - the history of accessed features and their related feature item IDs in the order the
- 15 features were visited (thus, a feature is added to the top of the accessed-feature history list when one of its feature items is accessed by - that is, presented to - the user whilst the feature is the currently visited feature).

If a visited-feature history list is kept, a history of accessed features can be embedded in it by providing each feature in the history with an associated flag to indicate whether or not

20 the feature was accessed whilst current. Although keeping a visited-feature history provides more information about the visit, it will inevitably use more memory resources than an accessed-feature history and in many cases it will only be desired to track features which the user has found sufficiently of interest to access an associated feature item. Where the purpose of the feature history is simply to keep a list of features (and related feature

25 items) that were of interest to the user, it may be desirable to exclude from the list features for which items were automatically presented but are not associated with exhibits (real or virtual) - that is, exclude features concerned with incidental information about the hall.

The feature history preferably covers the whole of the visit though it may alternatively only

30 cover the most recently visited/accessed features. In either case, the most recent several entries in the history list form what is hereinafter referred to as the "feature tail" of the user and provides useful information about the path being taken by the user.

The visit data held in memory 43 may further include details of a planned route being followed by the user, and a history of the locations visited by the user (this may be a full history or just the locations most recently visited - hereinafter termed the “location tail” of the user).

The service system 35 comprises the following main functional elements:

- A communications interface 50 for communicating with the mobile device 50 via the wireless LAN infrastructure 36.
- 10 - An internal LAN 51 (or other interconnect arrangement) for interconnecting the functional elements of the service system.
- A data store 52 for storing feature data and, in particular, a feature record for each feature with each record comprising the feature identifier, the subject of the feature, the corresponding real-world location and extent of the feature’s active zone, the IDs and media type of the or each associated feature item, and a flag which when set indicates that feature item presentation of an associated feature item is to be effected automatically upon delivery when the feature is being visited.
- 15 - A feature-item server 53 for serving an identified feature item to the mobile device 31 in response to a request from the latter.
- 20 - A location report manager 54 for receiving location reports from the location determination subsystem 40 of the mobile device and for passing on data from the reports to functional elements 55 and 56 (see below).
- A pheromone trail subsystem 55 for receiving location data, via manager 54, from all user mobile devices to build up trail data in a manner akin to the use of pheromones by ants.
- 25 - An item-data response subsystem 56 for receiving location and other data from the manager 54 in order to prepare and send a response back to the mobile device 31 that provided the location data, about what feature items it needs, or is likely to need, both now, in view of a feature currently being visited, and (where, as in the present embodiment, pre-emptive caching is implemented) in the near future. Subsystem 56 comprises a location-to-feature item translation unit 57 which can either be
- 30 implemented independently of the data held in store 52 or, preferably, be arranged to

operate by querying the store 52, the latter having associated functionality for responding to such queries. Subsystem 56 further comprises a prediction unit 58 for predicting, in any of a variety of ways to be described hereinafter, what feature items are most likely to be needed in the near future.

- 5 - A route planner 59 for responding to requests from the mobile device 31 for a route to follow to meet certain constraints supplied by the user (such as topics of interest, time available, person or tour to follow, an exhibit or facility to be visited, etc). In providing a planned route, the route planner will typically access data from one or both of the feature data store 52 and the pheromone trail subsystem 55. The route
10 planner 59 can conveniently hold a master map of the hall 10 for use by itself and the other elements of the service system 35, and for download to each mobile device 31 at the start of each new visit and/or whenever the master map is changed.

The functional elements of the service system 35 can be configured as a set of servers all connected to the LAN 51 or be arranged in any other suitable manner as will be apparent to
15 persons skilled.

The mobile device 31 and service system 35 provide a number of useful capabilities that will each be described in detail below after an overview of the general operation of the mobile device and service system during a visit. It is to be understood that the split of
20 functionality between the mobile device 31 and service subsystem 35 can be varied substantially from that indicated for the Figure 2 embodiment; indeed all functionality can be provided either entirely by the mobile device 31 (with all feature items being stored in the device) or by the service system 35 (with the presentation of feature items to a user being by means of fixed input/output devices located around the hall near the locations
25 associated with the virtual features).

In general terms, a user starting a visit can request a route to follow using the user interface 48 of the mobile device 31 to indicate parameters to be satisfied by the route. This route request is sent by the visit manager to route planner 50 and results in the download to the
30 mobile device 31 of a planned route. The path guide 49 then provides the user (typically, though not necessarily, only when asked) with guide indications to assist the user in following the planned route . Where the interface 48 includes a visual display, this can

conveniently be done by displaying a map showing the user's current location and the planned route; in contrast, where only an audio interface is available, this can be done by audio cues to indicate the direction to follow. A user need not request a planned route and in this case will receive no guide indications. A user may request a route plan at any stage
 5 of a visit (for example a route to an exhibit of interest).

As the user moves through the hall, the location determination subsystem 40 sends periodic location reports 62 (see Figure 3) to the location report manager 54 of the service system 35 via the wireless LAN 36. In addition to the user's current location, these reports
 10 typically include a user identifier (and possibly, additionally or alternatively, a type identifier indicative of any variable of interest such as, for example, the group of users to which the device user belongs or an activity being undertaken by the user), user/device profile data, and prediction-assist data for use by the prediction unit 58 in predicting what
 15 or more of the following: route data concerning any planned route being followed; the user's "location tail"; and the most recent feature (either the "most-recently visited" or "most-recently accessed") associated with the user, either provided alone or as part of the user's "feature tail".

20 When a location report 62 is received by the manager 54, it passes on the user's current location in the report to the pheromone trail subsystem 55 to enable the latter to build up trail data from all devices; additionally, the user and/or type identifier may be passed on to subsystem 55 if provided in the location report. The user's current location is also passed to the item-data response subsystem 56 together with any profile data and prediction-assist
 25 data in the location report 62. The item-data response subsystem 56 then constructs and sends a response 65 (see Figure 4) to the mobile device 31 that originated the location report.

More particularly, the location-item to feature translation unit 57 of subsystem 56 uses the
 30 data passed to subsystem to determine the feature, if any, currently being visited by the user and thus what feature items are relevant to the user in their current location. In doing this, the unit 57 may also use the supplied profile data to disregard both features that do not

relate to a subject of interest to the user and feature items of a media type that cannot be handled by the mobile device 31. The unit 57 may also use elements of the prediction-assist data (for example, the location or feature last encountered before the current one) to enable it to determine the direction of progression of the user and thus to select between
 5 feature items of a feature in dependence on the direction of approach of the user. This is done, for example, for the features associated with openings 25 in order to select a feature item appropriate to entering a room. The IDs of feature items identified by the unit 57 together with the identity of the corresponding feature and the status of the automatic presentation flag of the feature, form a first part 66 of the response 65 to be sent back to the
 10 mobile device 31. Where the current location does not correspond to the active zone of any feature, the first response part 66 simply indicates this.

A second part 67 of the item-data response 65 is produced by the prediction unit 58 and comprises a list of the feature items most likely to be needed in the near future by the
 15 mobile device 31; for each such feature item, the second response part 67 includes the feature ID, its type, size and probability of usage (discussed in detail hereinafter). Like the unit 57, the unit 58 uses supplied profile data to disregard feature items of features not of interest to the user or of a media type that cannot not be handled by the mobile device 31. The number of feature items identified in response part 67 is preferably limited (for
 20 example, to ten such items). The item-data response subsystem 56 then sends the response 65 back to the mobile device 31 of the user by using a return address supplied with the original location report 62 and passed to subsystem 56 by the manager 54.

Rather than having the prediction unit 58 provide a prediction each and every time the
 25 mobile device 31 sends a location report, it is possible to arrange for the prediction unit 58 only to operate when required by the mobile device 31 with the latter only requiring a prediction, for example, every nth location report or only after the user has moved a certain distance since the last prediction made by unit 58. Conveniently, the location report field used to carry the prediction-assist data is also used to indicate when a prediction is required
 30 by, for example, setting the field to a predetermined value when prediction is not required.

The item-data response received back at the mobile device 31 is processed by the visit

manager 47. If the first part 66 of the response identifies a feature (thereby indicating that the current location of the user corresponds to the active zone of feature), the manager 47 updates the 'current feature' data in memory 45 to the feature identifier and item IDs in the first response part. These item IDs are also passed to the cache manager 45 and are used by
 5 the latter to request immediate delivery of these items from the server 53 of the service system to cache 44, if not already present in the cache. If the feature history data held by memory 43 relates to visited, rather than accessed, features, and if the feature identifier and item IDs in the first response part 66 differ from the most recent entry in the feature history list, the latter is updated with the feature identifier and item IDs from the first response part
 10 66.

In the case that no feature is identified in the first part of the response 65, the 'current feature' data in memory 43 is set to null.

15 The manager 47 also determines whether the (first) feature item (if any) identified in the first response part 66 is to be immediately presented to the user, this determination taking account of the setting of the automatic presentation flag in the first part of the response, any user indication (stored, for example in the profile data) that all items are to be automatically presented, and any monitored indications of the user's interest in the
 20 currently-visited feature. Where a feature item identified in the first response part is to be immediately presented to the user, the manager 47 requests the item from the cache manager 45 (there may be a delay in the delivery of the item if it has not yet been received from the server 53). At the same time, if the feature history concerns accessed features the manager 47 updates the feature history with an entry corresponding to the feature identifier
 25 and item IDs forming the 'current feature' data; where the feature history although recording all visited features, provides for indicating whether a feature has been accessed, the manager updates the feature history accordingly.

With respect to the data contained in the second part 67 of the response 65, the visit
 30 manager simply passes this data to the cache manager 45 which determines whether or not to request from server 53 any of the items identified that are not already in the cache 44. The cache manger 47 in making this determination takes account of the probability that an

item will be needed in the near future and the available cache space. The cache manager 45 may decide to create additional cache space by flushing one or more items from the cache and/or by reducing the space they occupy, as will be more fully described hereinafter.

5

In this manner, the cache manager 45 seeks to ensure that the next item requested by the visit manager 47 as the user progresses to the next feature will already be in the cache 44.

Following the processing of an item-data response by the visit manager, whenever a feature item is accessed (presented) either as a result of the visit manager 47 determining that the current feature is of interest to the user or as result of the user specifically requesting the item (for example, after inspecting the list of items associated with the current feature), then where the feature history data records accessed feature information, the visit manager 47 checks if the feature associated with the accessed item is the current feature and, if so, updates the feature history to record the feature as an accessed one if not already so indicated.

10
15
20

The visit manager can also be arranged to keep a list in memory 43 of the individual feature items accessed.

Having described the general operation of the mobile device 31 and service system 35, a more detailed description will now be given of some of the functionality embodied in the arrangement of Figures 1 and 2..

25 Pheromone Trails

The location reports provided by the mobile device 31 and passed to the pheromone trail subsystem 55 serve as virtual markers in the virtual world corresponding to the hall environment. These markers are stored by the subsystem 55 and used to build up trail and other useful information about utilisation of the corresponding real-world environment.

30

In general terms (that is, without limitation to the specifics of the embodiment of Figures 1 and 2), the virtual markers left in whatever manner by a mobile-device user can be given a

variety of characteristics. For example, the markers can be arranged to reflect the nature of pheromones laid by social insects such as ants and have the following characteristics:

- the markers are left automatically;
- markers from different users are undifferentiated;
- 5 - markers have a value that diminishes both with time and with the distance from the point of marking;
- markers accumulate, that is the value of overlapping markers at a point is the sum of their values at that point;
- markers can be detected by all other users of mobile devices in the environment.

10

However, each of these characteristics represents a choice in some dimension and other choices are possible. For example:

- each marker may be left following a specific user action to do so in respect of that marker (that is, left deliberately);
- 15 - markers may be identified by their source;
- markers may be of different types to reflect different activities or intentions by the source;
- markers may be persistent;
- markers may be stored as distinct elements;
- 20 - perception of the markers may be limited to particular users.

Of course, a wide range of mixes of the above choices of characteristics (and of other characteristics) are possible and although the term "pheromone trail" is used herein to refer to the general arrangement of the deposition and use of virtual markers, this term should

25 not be taken as implying that any particular characteristic is being implemented in respect of the arrangement concerned or that the use of the markers is related to delimiting a trail . Furthermore, it is to be understood that implementation of any particular characteristic is not tied to either one of the mobile device 31 or service system 35 (indeed, the latter is not essential for the implementation of a pheromone trail arrangement where the devices can

30 communicate amongst themselves).

Whatever the detailed characteristics of the markers, the effect of their deposition as users

move around the physical environment is the generation of a marker "landscape" in the digital representation of that environment. The ridges, peaks, troughs and wells of this landscape reflect the number of markers laid in each part of the landscape and will typically (though not necessarily in all cases) also reflect the time elapsed since the markers were laid. The devices of other users are arranged to be able to sense this landscape enabling them to use various gradient and contour following applications to traverse it, for example to follow (or avoid) popular paths. In doing so, the intensity of marker accumulations can be indicated to users in a variety of ways; for example intensity levels can be represented through an audio signal whose loudness or frequency varied with that intensity or through a visual display.

Figure 5 depicts some of the implementation choices available when constructing an embodiment of the pheromone trail arrangement, these choices being arranged by processing stage according to a division of the pheromone trail process into five such stages (other divisions being possible). The five stages depicted in Figure 5 are marker deposition 80, storage 81, intrinsic behaviour 82 (applied to the stored data), application processing 83, and presentation. 84. These stages are carried out by corresponding functional blocks 85 to 89 depicted in Figure 6 with the storage block 86 having two sub blocks, namely a storage pre-processing block 90 and a memory block 91. Also shown, in dashed lines, in Figure 6 are the mobile device 31 and pheromone trail subsystem 55 of the Figure 2 embodiment positioned to indicate where the functional blocks 85 to 89 are disposed in that embodiment.

Considering first the marker deposition stage 80 (functional block 85), marker deposition can be done automatically, by deliberate user-action per marker, or by deliberate user confirmation of an automatically-generated series of latent markers representing a trail segment. Where markers are deposited (or generated) automatically, the frequency of deposition/ generation can be made time or distance dependent (with respect to the last marker) or can be arranged to occur at specific way points around the environment, for example, at virtual features (that is, when a user enters the active zone of the feature, with typically only one marker being deposited/generated per feature encounter). Depositing a marker when a feature is encountered does, of course, require the mapping between

location and features to have first been carried; this can be done either by arranging for this mapping to be effected in the user's mobile device or by arranging for the unit carrying out the mapping away from the device (for example, unit 57 in the Figure 2 embodiment) to deposit a marker on behalf of the device. However a marker is deposited/generated, each
 5 marker may have an associated user identifier and/or type indicator (indicating some special significance not specific to a user); each marker may also have a tag to indicate a desired decay behaviour – for example, where, by default, a marker is arranged to decay, a no-decay tag can be included which if set (or “true”) indicates that the marker concerned is not to be given the default behaviour of decaying in value with time.

10

The main choice presented at the storage stage 81 (functional block 86) is whether a marker is to be aggregated with previously stored markers deposited at the same location or stored as an individual marker along with any associated data. Whilst aggregation produces useful information, storing in an un-aggregated form has the advantage of preserving the
 15 maximum amount of raw data whilst leaving open the option to later on retrieve a copy of the data for aggregation; the disadvantages of not aggregating initially are the much greater storage capacity required and the delay in later on obtaining aggregated data. Where aggregation is effected, this can be done across all types of marker or for each type of marker separately. Typically aggregation is done by adding an initial strength value to the
 20 aggregated strength value already stored for the same “location cell” as that in which the marker was deposited where a location cell corresponds to a specific area of the real-world environment. Rather than a marker being allocated to one location cell only, the strength of the marker can be divided up between the nearest cells in proportion, for example, to the distance between the point of deposition of the marker and a center point of each of the
 25 nearest cells.

The intrinsic behaviour stage 82 (functional block 87) applies behaviours to the aggregated or non-aggregated markers such as reducing their strengths with time. Where a marker is individually stored, the marker can also be given a limited life determined as expired either
 30 when its strength falls below a certain level or when the marker reaches a certain age (for which purpose, a time-of-deposition time stamp can be stored along with the marker). Applying intrinsic behaviour is done, for example, by a process that regularly scans the

memory block 91, reviewing and modifying its contents as necessary. The intrinsic behaviour stage 82 may not be present either because no relevant behaviours are being implemented or because they are applied as part of the applications processes for using the stored data.

5

The application stage 83 (functional block 88) uses the stored data to carry out specific application tasks and may also apply behaviours, such as marker strength fall off with time, to the data copied from storage where such behaviours have not been applied earlier to the stored data itself. Typical application tasks include:

- 10 - where markers of all types are aggregated (either on storage or by the application), determining and following a “ridge” of the highest-strength marker aggregations corresponding to the most popular trail through the environment; a related application is one where a “trough” of the weakest (or zero) marker aggregations is followed;
- 15 - where markers are stored individually with user IDs and a strength fall-off with time behaviour has been applied to the stored data, following a trail left by a specific user, the strength of the relevant markers indicating the direction of movement along the trail;
- 20 - where markers are stored individually with user IDs and deposition timestamps enabling the trail laid down by each user to be followed, predicting or proposing a user’s future movement on the basis of the movement forward from that user’s current location of previous users whose trail leading to this location matches closely with the location tail of the subject user (that is, with the locations of the last several markers deposited by the current user);
- 25 - where markers are deposited on encountering a virtual feature and the markers are aggregated by type with a decay that is exponential in form with a time constant of half a day for example, determining the most popular features of a given type for the current day by determining the strongest aggregation of markers of that given type.
- 30 As regards the presentation stage 84 (functional block 89), how the output of an application is presented to a user will depend on the nature of that output and the interface modalities available. Typically, where an application task has determined a trail to follow or the most

popular features, this can be presented to the user on a map display whilst if an application is arranged to provide real time guidance along a path, this may be best done using audio prompts.

- 5 Implementation details of the functional blocks 85-89 for any particular combination of the available choices discussed above will be apparent to persons skilled in the art.

It should be noted that multiple combinations of choices can exist together; for example, markers can be arranged to be deposited by a mobile device both at fixed time intervals and
 10 every time a feature is encountered and a marker can be both aggregated upon storage as well as an individual copy being kept.

With respect to the embodiment of Figures 1 and 2, the pheromone trail subsystem 55 is arranged to store markers of three different types, namely:

- 15 - **“tour” markers** deposited in the form of location reports 62 by a tour guide and serving to delineate a proposed route around the hall. These markers are each deposited by deliberate act of the tour guide and have an associated “no-decay” tag as well as an indicator of their type. Preferably the type indicator has an associated sub-type that identifies a specific tour. Since each specific tour will have only one set of
 20 markers associated with it, storing the tour markers on the basis of aggregating markers of the same type and sub-type deposited in the same location is the same as storing the markers individually and either approach may be adopted. The stored markers are not decayed with time.
- **“normal” markers** deposited in the form of location reports 62 by the mobile
 25 devices 31 of visitors, these markers being deposited at fixed time intervals and being subject to aggregation on storage with other markers of this type deposited in the same location cell (that is, an initial strength value associated with a newly deposited marker is added to the aggregated strength value associated with the marker aggregation for the cell in which the new marker has been deposited). The
 30 strengths of the marker aggregations are decayed with time but over a long time period. These aggregated “normal” markers serve to indicate the most popular trails, reflecting both the number of users traversing these trails and the time spent on them.

- **“feature” markers** deposited by the unit 57 each time it determines from data in a location report that the device sending the report is in the active zone of a feature. If, as ifs preferred, the prediction assist data in the location report contains current feature data, then deposition of a feature marker can be restricted to when a user first enters the active zone of the feature, this being achieved by comparing the identity of the current feature as determined by unit 57 with the current feature noted in the location report and only depositing a marker if the two differ. The feature markers are aggregated in feature cells held by the unit 55 and are decayed over a period of an hour to give a picture of the current popularity of the features. Feature cells are simply location cells covering an area corresponding to the active zone of a feature.

The stored markers are put to use for route planning / following, feature popularity review, and prediction purposes. With respect to route planning, when the visit manager 47 of a mobile device 31 requests a route from the route planner 59 of the service system, the latter can ask the application task block 88 of the pheromone trail subsystem 55 to access the stored marker data and propose a possible route based either on the tour markers or the aggregated normal markers. Thus, the route planner, where provided with a subject of interest to the user by the visit manager 47, can be arranged to map this subject to a particular tour sub-type and then retrieve the set of locations of the corresponding tour markers stored by the subsystem 55; these locations are then used to provide a route plan back to the mobile device 31. As described above, no sequence information is stored with the tour markers and whilst this will generally not be an issue, it is possible to provide for the tour markers to carry sequence information in a number of ways, the simplest of which is to associate a sequence number with each tour marker as it is deposited, this number being incremented for each successive marker and being stored along with the marker. An equivalent way of providing sequence information is to incrementally increase/decrease the strength value assigned to each marker as it is deposited; since the tour marker do not decay, this strength value remains and effectively serves as a sequence number indicating the direction of progression of the tour.

The route planner 59 can be arranged to request the subsystem 55 for the most popular route around the hall 10 as indicated by ridges of higher-strength accumulations of normal

markers, or for the least crowded route as indicated by troughs of zero or low-strength accumulations of the normal markers. Of course, the route planner 59 will typically have been requested by a user to provide a route that takes the user to features relating to a particular subject or even to a set of user-selected features; if the route planner decides that there is no relevant pre-planned tour it can use, or if the user has specifically asked for a popular or a least crowded route, then the route planner will use the normal-marker aggregations to aid it in planning a route between the selected features. This can be done by first selecting an order in which to visit the features and then asking the application task block 88 to provide the most popular / least crowded route between each successive pairing of features in the order they are to be visited. Alternatively, the actual order of visiting of the features, as well as the route between each feature, can be determined according to the peaks and troughs of the accumulated normal marker landscape, preferably with account being taken of the total distance to be traveled by the user. In this case, since the application task block 88 has more immediate access to the stored marker accumulations, it may be appropriate for the route planner to hand over the whole task of planning a route to the task block 88.

Rather than determining a route by following ridges or troughs in the accumulated-marker landscape, the route planner can be arranged to determine a route by avoiding ridges or troughs.

Another possible usage of the pheromone trail subsystem 55 in respect of providing route information involves the deposition by a first user of user-specific markers that are not aggregated but are arranged to decay in strength over a period of an hour or so. These markers would enable a second user to request the route taken by the first user (for example, by means of a request sent from the visit manager 47 of the second user's mobile device to the route planner 59), the markers deposited by the first user then being accessed to determine the route taken by the first user and their direction of progression as indicated by the current strengths of the markers. This service (suitable for a parent wanting to track a child) can be made private with only the users involved being able to access the relevant marker data and can be provided as a fee-based service.

A similar type of usage involves all members of a group having markers of a type specific to that group, the markers being aggregated on storage. This would enable an overview to be obtained of what the group did during a visit and if the markers concerned did not decay (though typically given a lifespan limited to the day of the visit) and were deposited
 5 at fixed time intervals, it would also enable the popularity of different visited features to be discerned.

Although in the foregoing examples of the use of the pheromone trail system in the embodiment of Figures 1 and 2, the route information derived from the stored markers has
 10 been passed back to the mobile device for storage in the visit data memory 43 as a route to be followed, it is also possible to have a more dynamic interaction between the mobile device and the stored marker data. Thus, for example, the mobile device 31 can be arranged to query the pheromone trail subsystem 55 continually as to the next location to move to in order to follow a ridge or trough of the marker landscape or to follow a trail laid down by a
 15 specific user.

With regard to the use of the deposited marker data for feature popularity review, if a user wishes to ascertain the current relative popularity of the features (or, in user terms, of the exhibits with which the features are associated), the user causes the visit manager 47 to
 20 send a request to the pheromone trail subsystem 55. The task block 88 of the subsystem 55 then accesses the feature marker accumulations of the feature cells and uses the strengths of these accumulations to determine the current relative popularity of the features. This popularity data is then returned to the requesting mobile device for presentation to the user. If a longer term view of the popularity of the features is required, then the task block 88
 25 accesses the normal marker aggregations for the feature cells, these aggregations having a longer decay period and, unlike the feature marker accumulations, having a strength that reflects the dwell time at each feature as well as the number of visits.

In respect of use of the deposited marker data for prediction purposes, this involves using
 30 the current location or location trail of a user to make predictions as to where the user is likely to go next having regard to what others have done as indicated by the relative strengths of the accumulations of normal markers in location cells adjacent the one in

which the user is currently located. If location tail data is available, the strengths of marker accumulations in location cells just visited by the user (and possibly also of the cells on either side of such cells) can be scaled down to reflect the fact that the user is less likely to visit those cells; however, if the geography of the hall or the layout of features of interest to the user is likely to cause the user to turn around, then such scaling down is not effected. Making predictions of the user's future path in this manner is carried out by the application task block 88 of the pheromone trail subsystem. As will be further described below, this future path prediction capability can be used by the prediction unit 58 to determine what feature items are likely to be needed in the near future.

It will be appreciated that many other applications are possible for the pheromone trail arrangements discussed above.

15 Path Following

However a route is determined by the route planner 59, whether by using the pheromone trail subsystem 55 or in some other manner, details of the planned route are passed back to the mobile device 31 for storage in the memory 43. Alternatively, a route to follow may have been determined in the device itself, for example by the user specifying on the stored map locations to be visited and the visit manager 47 locally determining the shortest path between these locations. Typically, the route will have been specified by a series of locations defining a path. The path guide unit 49 is operative to use these stored details to provide guidance to the user for following the path concerned. Whilst the path guide unit 49 can be arranged to use a visual display of user interface 48 to provide this guidance, two embodiments of unit 49 are described below for using non-verbal audio output to guide a user along a specified path (referred to below as the "target" path).

In the first embodiment of unit 49, stereo audio cues are provided to the user that indicate how close the user is to the centreline of the target path (see chain-dashed line 100 in Figure 7). These cues are typically presented through stereo headphones 60 but other audio output arrangements are possible such as small shoulder-mounted loudspeakers. In Figure 7 the actual path taken by the user is indicated by arrow 101 with the user's current

location being at the point referenced 102. At point 102, arrows 103 and 104 respectively indicate the direction of pointing of the path centre-line 100 and the current direction of moving of the user, the angle between the two having a value of X degrees. The perpendicular distance from the target path 100 to the user's current location 102 has value "d".

In general terms, the first embodiment of the unit 49 operates by repeatedly carrying out the following steps:

- determine the distance "d" using the most recent location data stored in memory 43 by the location determination subsystem 40 and the details of the target path, also stored in memory 43;
- render audio cues in each of the stereo channels according to a function that relates the distance "d" to an audio characteristic such as frequency or intensity.

The cues to the left and right audio channels are preferably controlled to vary in a complementary manner – thus as the value of "d" changes, the audio characteristic used to represent the distance "d" is changed in one sense for the left channel and in the opposite sense for the right channel. For example, and as illustrated in Figure 8, both channels can be arranged to carry a tone of the same frequency when the user is on the centreline of the target path. As the user moves away from the centreline, the frequency of the tone in one channel is caused to rise, and that in the other channel is caused to fall. In order to move back to the centreline, the user must move to bring the left and right channel tones back to a common frequency. Such balancing of the frequency of two tones is akin to tuning a musical instrument and can exploit the familiar phenomenon of minimising the beat frequency between the tones.

In the Figure 8 example, as a user moves off the centreline to the left, the frequency of the tone carried by the left channel increases and that of the right channel decreases; it will be appreciated that the reverse policy could be applied - that is, the frequency of the left-channel tone could decrease as the user moves to the left whilst that of the right-channel tone increases.

It will also be seen in the Figure 8 example that changing the frequencies of the channel tones is limited to positive and negative values of "d" below a certain limit magnitude "D"

- outside this range of values, the channel tone frequencies are both zero. In effect, the threshold value D defines a corridor of width 2D centred on the centreline of the target path

5 100. In Figure 7, the boundaries of this corridor (the "walls" of the corridor) are indicated by dashed lines 106, 107.

The complementary variation of the controlled audio characteristic (frequency, intensity, etc.) of the left and right audio channels need not extend across the whole range of the

10 parameter "d" for which variations of this characteristic are produced. Thus in the Figure 8 example, as the user moves to the left of the centreline the frequency of the right-channel tone can fall to zero prior to the left-channel tone reaching its maximum frequency at distance "-D" from the centreline. It is also possible to arrange for there to be a plateau region either side of the target-path centreline within which the tones in the two channels

15 do not change as the user moves towards and away from the centreline. It is further possible, though not preferred, to do away with any range in which complementary variation takes place - for example, in a central region either side of the target-path centreline no tones are generated but as the user moves outside of this region to one side or the other a tone is produced in the audio channel on that side.

20

Furthermore, the unit 49 can be arranged to vary a second characteristic of the audio signal to indicate the desired direction of travel along the path. For example, where channel tone frequency is used to indicate the distance "d" from the centreline as described, the loudness (intensity) of the tones can be increased progressively along the path from start to finish (or

25 from the start of a path segment to the end of the segment) to indicate the direction to move along the path. A drawback of this approach is that it requires the user to be moving before it is apparent whether or not the user is moving in the correct direction. If a user stops along the path and turns for any reason, it will not be immediately apparent in which direction the user should set off when wishing to re-commence following the target path.

This drawback is overcome in the specific implementation of the first embodiment of the path guide unit 49 described below with reference to Figures 9 and 10. In this implementation, unit 49 has four main states (see Figure 9), namely:

- a STANDBY state 110 in which unit 49 resides when not actively providing audio cues to the user;
- a START state 111 entered when the unit first starts to give audio cues to the user for moving along a target path – in this state, the user is always initially located on the path centreline and is assumed not to be moving;
- a MOVING state 112 in which the unit resides when operating to give audio cues to a user that has been detected as moving; and
- a STOPPED state 113 in which the unit resides when operating to give audio cues to a user that has been detected as having stopped.

In the Figure 9 state diagram, conditions to be satisfied to transit between states are indicated by the legends in square brackets on the arcs between states, and actions associated with transitions are indicated by legends in round brackets.

In each of the active states (that is, all states except the Standby state 110), the unit 49 operates according one of three control regimes A, B and C in dependence on the angle Y° (see Figure 10) between the current direction of pointing 103 of the target-path centreline and a control direction 117. More particularly, control regime A applies when the angle Y has a magnitude of less than α° , control regime B applies when the angle Y has a value in the range $(+\alpha^\circ \text{ to } +\beta^\circ)$ or $(-\alpha^\circ \text{ to } -\beta^\circ)$, and control regime applies for all other values of angle Y .

For the Start state 111 and the Stopped state 113, the control direction 117 is the current direction of facing of the user. This direction is measured, for example, by an electronic compass mounted on stereo headphones of the user to give an absolute direction of facing of the user, the direction of pointing of the centreline of the target path also being known in absolute terms through an absolute reference direction stored for the map of the hall held in memory 43 (the target path being specified with respect to the map).

For the Moving state 112, the control direction 117 is the current direction of moving of the user (that is, direction 104 in Figure 7 in which case angle Y equals angle X of Figure 7). The direction of moving of the user is determined, for example, by using the last two measured locations of the user.

5

The behaviour exhibited by the unit 49 in each of the control regimes is as follows:

- Regime A - the user is presented with stereo audio cues with each stereo channel having a characteristic that varies with the distance "d" as described above, for example, with respect to Figures 7 and 8;
- 10 Regime B - as for regime A with the addition of a world-stabilized sound cue generated so as to appear to come from a point in space a long way off along a projection of the current direction of pointing of the target-path centreline. By "world-stabilized" is meant that the direction from the user's current position to the sound-cue source point does not vary relative to a real-world fixed reference direction as the user turns their head or body. Generation of
15 a world-stabilized sound cue can be readily effected using a 3D audio spatialisation processor provided that the centreline of the audio output arrangement is known (where the audio output arrangement is stereo headphones, this centreline is simply the direction of facing of the user). This sound cue is generated to be a sound distinct from the audio cues used
20 to indicate the distance "d" and serves to indicate to the user the direction in which the user should face to be pointing along the target path; this sound cue is referred to as the target-path direction cue below.
- Regime C - Only the target-path direction cue is provided to the user to enable them to
25 correctly orientate themselves for moving along the target path.

In operation, the user first selects a target path whilst the unit 49 is in its Standby state 110, this target path starting from the user's current location. Upon a user command being given to start the guidance process, the unit 49 transits to the Start state 111 in which the user is
30 positioned on the centreline of the target path and the control direction 117 is the user's direction of facing. If the user is facing within α° of the centerline direction of the target path, the unit 40 operates in control regime A to provide the user with stereo audio cues

that will initially be balanced and the user can set off in their direction of facing knowing they are proceeding in roughly the correct direction. If, however, the user is facing away from the target-path centerline direction of pointing by more than α° , unit 49 will operate in control regime B or C providing the user with the target-path direction cue to permit the user to turn in the correct direction before starting off. By also providing the stereo audio cues in regime B when the user's initial direction of facing is not completely inappropriate, the user can start moving without having to fully orientate themselves first.

As the user moves off, the fact that they are moving is detected (generally by detecting the a change in sensed location) and the unit 49 transits to the Moving state 112. Although the unit 49 continues to operate in control regime A, B or C according to the value of the angle α° , the control direction is now the direction of movement of the user – in other words, the direction of facing of the user becomes irrelevant to the cues provided to the user and they can turn their head from side to side as they move along without this affecting the cues they receive. If the user starts moving in a direction that is greater than α° (but less than β°) away from the path centreline direction, the user will not only hear changes in the stereo audio cues caused by changes in the distance “d”, but also hear the target-path direction cue to help them correct their direction of movement.

If the user stops for more than a short period (as detected, for example, by the detected location of the user remaining unchanged for n readings where n is an integer greater than one), the unit 49 transits to the Stopped state 113. This state is similar to the Start state 111 except that the user is not necessarily positioned on the path centreline. Thus, the user will receive audio cues regarding the distance “d” and/or the path centreline direction according to the angle between the latter and the direction of facing of the user.

Upon the user moving off again from the Stopped state 111, the unit 49 transits back to the Moving state 112.

Upon the user arriving at the endpoint of the target path, the unit 49 transits out of the Moving state 112 back to the Standby state 110 and, in doing so, provides an appropriate indication to the user that they have arrived at their destination.

- If whilst in the Moving state 112, a user should move outside of the corridor defined by the value D for the distance “d”, the unit 49 temporarily transits back to the Start state 111 and as it does this, the target path is re-aligned to pass through the user’s current location. In other words, the path is moved to put the user back on it. Since the user is moving, the unit 49 quickly returns to the Moving state 112. A similar path re-alignment by transiting the unit 49 back to the Start state 111 is also effected if the user remains in the Stopped state 113 for more than a predetermined timeout period.
- 10 Finally with respect to the Figure 9 state diagram, the user can by user command cause the unit to move from any of the active states 111-113 back to the Standby state 110 in which no audio cues are given; the user may also, by user command, cause the unit 49 to transit from the Moving state or Stopped state to the Start state 111 to cause the target path to re-align to have its centreline pass through the user’s current location.
- 15 Many variants are, of course, possible to the above-described implementation of the first embodiment of unit 49. For example, in control regime B rather than providing a target-path direction cue of the form described to supplement the stereo audio cues indicative of the distance “d”, these latter cues can simply be supplemented by a distinctive sound in the
- 20 audio channel of the side of the target path to which the control direction is pointing. In this case, regime C can either be this distinctive sound alone or the target-path direction cue as before. A further alternative to using the target-path direction cue in control regime B is to make the stereo audio cues that are indicative of the distance “d”, also responsive to the fact that the control direction is pointing towards one side of the path and away from
- 25 the other. This can be done, for example, by notionally displacing the location of the user perpendicularly to the path centreline towards the side of the path pointed to by the control direction, the amount of this displacement being dependent on the magnitude of the angle Y° ; the value of the distance “d”, and thus the sound of the stereo audio cues, is then determined based on the notional position of the user after the notional displacement.
- 30 In a further variant, the value of α° or β° is set to zero – in other words, control regime A or control regime B is eliminated. With regard to the determination of the position of the

user relative to the target path, whilst in the above-described arrangements this has involved a direct measure of the perpendicular distance “d” from the target-path centreline to the user’s current position, it will be appreciated that indirect measures are alternatively possible such as determining the distance of the user from the nearest corridor wall 106,
 5 107.

In a second embodiment of the path guide unit 49, a 3D audio spatialisation processor is used to project a virtual audio beacon ahead of the user along the target path. As the user
 10 approaches (or arrives at) the location of the virtual beacon, the beacon is re-positioned further along the path. This process repeats until the user has traversed the entire path. The process is illustrated in Figure 11 in which the target path is indicated by chain-dashed line 120 passing from start point 121 to end point 122, and a partial piecewise linear approximation of the target path is indicated by lines 123A, 123B and 123C; as will be
 15 seen from the following description, the user is guided to follow this piecewise linear approximation rather than the exact target path.

When the user is positioned at start point 121 and the unit activated for guiding the user along target path 120, the unit determines at least a first segment 123A of the piecewise
 20 linear approximation to the target, this approximation being generated according to a heuristic which, for example, both keeps the area between this first segment and the target path to below a predetermined limit value, and keeps the length of the segment to no more than a predetermined maximum length. After determining this first segment 123A, the unit 49 determines a position 125A for a virtual sound beacon such that it lies a fixed distance
 25 “K” beyond the end of the first segment 123A in the direction of extent of the latter. The unit 49 then uses its 3D audio spatialisation processor to produce a world-stabilised virtual sound beacon at this position 125A in the sound field of the user, the output of the 3D audio spatialisation processor being via stereo headphones 60 or other suitable audio output devices (such as the shoulder mounted speakers previously mentioned). Also as
 30 previously mentioned, in order to render the virtual beacon in a world stabilized position, the unit 49 is provided with the direction of facing of the user’s head or body, depending on where the audio output devices are carried. An electronic compass mounted with the

audio output devices can be used to provide the required direction of facing data to the unit 49.

As a result, regardless of the direction of facing of the user, the user is provided with a
5 sound beacon positioned in space to indicated the direction in which the user should move
to follow the target path 122.

The user now sets off towards the position 125A of the virtual sound beacon. Upon the
user approaching to within distance "K" of the sound beacon, the above process is repeated
10 with the user's current position being taken as the start of the target path (whether or not
actually on the target path). Thus, a second linear piecewise approximation segment 123B
is determined and the virtual sound beacon is re-positioned to appear to be at a location
125B a distance "K" beyond this newly-determined segment in the direction of extent of
the latter.

15 In this manner, the virtual sound beacon is moved in a succession of steps to guide the user
along a piecewise linear approximation of the target path until the user reaches the path
end point 122.

20 It should be noted that where the target path includes a long straight section, this will be
split up into several segments by the above process in order to keep the segment level
down to the aforesaid predetermined maximum length thereby ensuring that the virtual
sound beacon is never more than that distance plus the value "K" beyond the user. This is
illustrated in Figure 11 by the third segment 123C for which the virtual beacon has been
25 located at a position 125C which is only part way along a straight section of the target path.

Many variants are possible to the above-described second embodiment of the path guide
unit 49. For example, the distance "K" may have a value of zero.

30 Furthermore, rather than providing only a single virtual sound beacon, one or more further
beacons can be provided beyond the first beacon, for example, at positions at the end of (or
a distance "K" beyond) second and third segments of the piecewise linear approximation

of the target path – in this case, the maximum length of each path segment will typically be less (for example, half) of that used in the case where only a single virtual beacon is being presented. For example, and as illustrated in Figure 12 for target path 120, three virtual sound beacons can be used each positioned at the end of a corresponding piecewise linear approximation segment, the first three of these segments 126A, B and C being shown with the beacons located at positions 127A, 127B and 127C corresponding to the end of respective ones of these segments.

The beacons can be caused to vary in intensity, frequency, or some other audible characteristic, to indicate the order in which they should be approached with each beacon sounding in turn (potentially with overlap) in a cyclic manner with the further sounds being quieter. In the Figure 12 example, graph 129 indicates the variation of sound volume v with time t for each of the first, second and third beacons (the first beacon being the one nearest the user along the target path and third beacon being the one furthest away from the user along the target path). The sound frequency (or other audible characteristic) can be varied in conjunction with changes with the volume of each beacon to represent distance between the beacon and the user, the sound frequency decreasing, for example, as the user approaches the beacons.

The virtual beacons are rendered to appear static and as the user approaches or reaches the first, it is removed and replaced by a new distant virtual sound beacon, last in the series of three beacons, positioned at the end of a further piecewise linear approximation segment. The new beacon may be caused to appear just before, at the same time as, or just after the first beacon is caused to disappear. This process of replacing beacons as they are approached or reached is repeated as the user moves along the target path.

As a further general variant, it is possible, though not preferred, to arrange for the beacon or beacons to remain a constant distance ahead of the user, at least over a substantial portion of the target path, as the latter seeks to move towards them.

Caching of Feature Items

As described above, the mobile device 31 is arranged to pre-emptively cache feature items in cache 44 in dependence on their respective probabilities of being required in the near future; these probabilities being determined by the prediction unit 58 of the service system
5 35 using information (in particular, the prediction-assist data) provided in the location reports from the mobile device 31. In this manner, the latency inherent in fetching feature items from the feature item server 53 only when needed is avoided.

The prediction unit 58 can operate on the basis of any one or more of a variety of different
10 techniques for predicting which feature items will be needed in the near future. A number of these techniques are described below, these techniques being divided into two groups, namely a first group A covering techniques that do not use visit data concerning previous users, and a second group B that rely on such previous-users visit data.

15 It should be noted that in the following the probability of an item being needed (also referred to as the probability of usage of the item), is used to encompass both the probability of an item being definitively requested (that is, not on a probabilistic basis) for delivery to the mobile device and the probability of an item being accessed at the device by the user. The fact that these two probabilities are different in the Figure 2 embodiment is
20 because the service system and mobile device operate on the basis that all items associated with a currently visited feature are downloaded into the cache 44 of the device, regardless of their probability of being accessed by the user. The probability that a particular item will be requested for delivery to the device is thus the same as the probability that the user will visit the associated feature. Had the service system and mobile device simply been
25 arranged to non-probabilistically request delivery of an item only when accessed by the user, the probability of an item being requested for delivery would be the same as the probability of that item being accessed. Notwithstanding the fact that in the Figure 2 embodiment all items associated with a current feature item are requested for delivery, prediction of what items may be needed in the near future need not be restricted to use of
30 the probability of a feature item being non-probabilistically requested (as indicated, for example, by the probability of the associated feature being visited), and can alternatively be based on the probability of the user accessing a particular item (or of accessing at least one

of the items associated with a feature, all these items then being considered as having the same probability of access). Consolidating the foregoing, the probability of usage of an item can be based on the probability of a feature being visited or accessed, or of a feature item being accessed by the user or non-probabilistically requested for delivery.

5

Furthermore, regardless of the prediction technique being used, the prediction unit 58 may, as already mentioned, filter out from its prediction process all feature items that do not relate to a subject of interest to the current user or are of a media type incompatible with the mobile device of that user. In fact, rather than filtering out all feature items concerning

10 subjects in which the user has not expressed interest, the probabilities associated with these items regarding their likely use in the near future can be appropriately adjusted to take account of the user's apparent lack of interest in them.

A – Prediction not based on Visit Data Concerning Previous Users

15

1. ADJACENCY OF FEATURES TO CURRENT LOCATION (OPTIONALLY WEIGHTED AGAINST WAKE FEATURES)

This is the simplest prediction technique and in its basic form takes the current location of the user and determines the closest features (typically by reference to the data held in the feature data store). The probability of usage of the feature items associated with these features is based on the probability of the features being visited and is thus set to fall off in dependence on the distance of the feature concerned from the user's current location. For example, all features within a 30 meter radius of the user's current location are determined and the probability of usage of an item

20 associated with a feature r meters away is set to:

25

$$(30-r)/30$$

This basic technique can be modified to reduce the probabilities of usage of feature items associated with features that the user has recently passed (and is therefore less likely to visit in the immediate future). These features, referred to below as "wake" features, are identified by the prediction unit 58 using location history data of the

30 current user – in present embodiment this data is supplied in the form of the user's location tail provided as part of the prediction assist data. As depicted in Figure 13,

the immediately preceding location 130 (or locations) of the user are used, together with the user's current location 131 to determine a "wake" zone 133; the probability of usage of any feature item associated with a feature 134 lying in the wake zone 133 (that is, a wake feature) is then weighted by a factor between 0 and 1. It will be appreciated that the wake zone 133 could be divided into sub-zones each having a different associated weighting factor according to a perceived reduced probability of usage of feature items for features in such sub-zones.

2. ADJACENCY OF FEATURES TO PLANNED ROUTE

If the user is following a planned route and data about the next portion of this route is included in the prediction assist data, the prediction unit 58 can use the route forward of the user's current location to determine the features next to be encountered along the route (either on the route or adjacent to it). The probability P of usage of a feature item associated with such features is again based on the probability of a feature being visited and is set according to both the distance l of the feature along the forward route (the greater the distance, the lower the probability) and the perpendicular distance d of the feature off the route (again, the greater the distance the lower the probability but this time the fall-off rate is much faster than for distance along the forward route). Figure 14 illustrate, by way of example, a linear fall off of probability with distances l and d giving a defining plane 139.

It may be noted that where the route being followed is a standard route for which route data is held by the route planner 59, then the route data included in the prediction assist data can simply be an identifier of the route, the prediction unit using this identifier to retrieve the route details from the route planner 59. It may also be noted that a planned route may be defined in terms of features to be visited rather than as a path; in this case, the probability of usage of feature items for features on the route is set simply by their order from the current point onwards; other features not on the route can still be included in the prediction according to their adjacency to the features on the route (or to a direct path between them). It may be further noted that having a planned route stored in visit memory 43 is not necessarily to be taken as a sufficient condition that the user is following a planned route; one or more

additional conditions may be required such as, for example, the user is actively using the path guide unit 49 to follow the planned route, or the last two/three features visited have all been on the planned route. The determination as to whether a planned route is being followed is preferably made in the mobile device 31.

5

3. ADJACENCY OF FEATURES TO FUTURE TRACK PREDICTED FROM MOVEMENT HISTORY

Where the user's recent movement history is available to the prediction unit 58 (for example, as a result of the user's location tail being included in the prediction data), then the unit 58 can use this information to predict the user's track in the immediate
 10 future. Thus, if the user's location tail is available to the prediction unit 58, a smooth curve passing through the locations in this tail can be determined and continued to predict the user's future track. This track can then be used in much the same manner as a planned route as described above, that is, the features lying on or near the track are identified and the probability of usage of feature items associated with these
 15 features is set in dependence both on the distance of the features concerned along the track and on their distance off the track (c.f. Figure 14).

20

Rather than predicting the user's future track on the basis of their location tail, this track can be predicted from a knowledge of the user's current location and direction of moving as determined for example, by the direction of facing of the user's body as measured by an electronic compass carried by the latter.

B – Prediction based on Visit Data Concerning Previous Users

25 This group of prediction techniques use visit data concerning previous users. This visit data can be collected in any suitable manner. For example, the visit data can be obtained by storing in a mobile device 31 during a visit, time-ordered lists of all locations and features visited, and all feature items accessed and where they were accessed. At the end of the visit, the stored data is uploaded to the service system for organization and use by the
 30 prediction unit 58. It is alternatively possible to arrange for the visit data to be collected by the service system as a user progresses through a visit. Furthermore, where prediction is based on location / feature trail information, the pheromone trail subsystem 55 can be used

to provide the required visit data.

4(a). SAME LOCATION – TRACK PREDICTION - FEATURE PREDICTION (OPTIONALLY WEIGHTED AGAINST WAKE FEATURES)

5 This prediction technique simply uses the user's current location to predict where the user is likely to go next on the basis of where previous users have gone from this location (the prediction unit 58 may, for example, query the pheromone trail subsystem for such information). Given the most likely future track(s) of the user, the features that will be encountered along or near the track are determined followed by
10 the probability of usage of the associated feature items; this is effected, for example, in a manner akin to that used in prediction technique (3) above. If more than one future track is considered, the probability of use of each track is used as an additional weighting factor for the probability of usage of the feature items. A weighting can also be introduced to reduce the probability of usage of feature items associated with
15 wake features as described above with reference to prediction technique (1).

4(b). SAME LOCATION – DIRECT PREDICTION OF FEATURE/ FEATURE ITEM (OPTIONALLY WEIGHTED AGAINST WAKE FEATURES)

20 Rather than predicting feature/feature item usage indirectly by first predicting a future track for the user based on the tracks taken by previous user's from the current location, it is possible to take the user's current location and use it to predict directly from the previous-users visit data what features will probably be visited or accessed next - or even more directly, what feature items are likely to be visited / delivered to the mobile device / accessed in the near future. This is done by organizing the
25 previous-users visit data on the basis of what features are most commonly next visited or accessed or what feature items are next delivered to or next accessed by users who have been at the current location. Again, a weighting can be introduced to reduce the probability of usage of feature items associated with wake features.

30 5(a). SAME RECENT MOVEMENT HISTORY - TRACK PREDICTION - FEATURE PREDICTION

This technique is similar to prediction technique (4a) but makes its future track prediction based on where previous users with the same recent movement history

(typically, with the same location tail) have gone from the current location. Of course, since previous locations visited by the user are inherently taken into account by this technique, it is inappropriate to adversely weight the usage probabilities of items associated with wake features as was optionally done for prediction technique (4). It should be noted that in order to be able to identify previous users with the same recent movement history, the movement data of previous users needs to be available in an un-aggregated form.

5(b). SAME RECENT MOVEMENT HISTORY - DIRECT PREDICTION OF FEATURE/ FEATURE ITEM

This technique is similar to prediction technique (4b) but uses visit data from previous users with the same recent movement history (typically, with the same location tail) in order to determine what features are likely to be visited or accessed in the near future, or what feature items are likely to be visited / delivered to the mobile device / accessed in the near future.

6. SAME MOST-RECENTLY VISITED FEATURE - PREDICTION OF FEATURE/ FEATURE ITEM (OPTIONALLY WEIGHTED AGAINST RECENTLY-VISITED FEATURES)

This prediction technique does not use location data but bases itself on visited feature data. More particularly, the prediction unit 58 uses the current feature identified by the location-to-feature translation unit 57 or, if the user's current location does not correspond to a feature, the user's most-recently visited feature as identified in the prediction assist data included in the latest location report. Given this feature, the prediction unit 58 accesses a stored table 140 (see Figure 15) which for each feature $F1$ to FN keeps a count of the next feature visited by each previous user. These count values enable the unit 58 to determine the probability of each feature being the next feature visited, these probabilities then being applied to the feature items associated with each feature as the probability of usage of those items. If the prediction assist data includes the feature tail of the user, this can be used to reduce the probabilities associated with the features in the user's feature tail.

The table 140 lends itself to dynamic updating since if the unit 57 identifies a feature - for example feature $F(N-3)$ - that is different to the most-recently visited feature -

for example, feature F5 - identified in the prediction assist data, this indicates that the user has moved from the feature F5 to the feature $F(N-3)$ so that the count value in the table cell at the intersection of row F5 and column $F(N-3)$ should be incremented to reflect this.

5

It should be noted that a single access to the table 140 will only give probabilities regarding the next feature to be visited. However, it is possible to look further ahead by accessing the table again in respect of the most-probable next feature (or features) in order to derive probabilities in respect of the next-but-one feature to be visited. By repeating this process, a forward-looking probability graph can be built up to any required depth.

10

It should also be noted that it is possible to provide table 140 with the next-visited feature probability data replaced by next-accessed feature data where, as explained above, an accessed feature is a feature having at least one associated item that has been accessed – presented to – the user. Alternatively, the next-visited feature count data in table 140 can be replaced by probability data about the next feature item visited / requested for delivery (or delivered) to / accessed by, the user after the current/most-recent feature visited.

20

7. SAME MOST-RECENTLY VISITED FEATURE HISTORY - PREDICTION OF FEATURE/ FEATURE ITEM

This prediction technique matches the user's recent visited-feature history (their visited-feature tail) to the visited-feature histories of previous users visiting the user's current or most-recently visited feature. Having identified previous users with matching feature tails, the prediction unit 58 analyses the visit data of these previous users to determine the probabilities associated with the user next visiting the other features and thus the probability of usage of the associated feature items. As with prediction technique (6), rather than predicting the next feature to be visited, the previous visit data can be organized to enable the unit 58 to predict the next accessed feature (and thus feature items likely to be needed) or the next feature item visited / requested for delivery (or delivered) / accessed.

25

30

8. SAME MOST-RECENTLY ACCESSED FEATURE - PREDICTION OF FEATURE/ FEATURE ITEM
(OPTIONALLY WEIGHTED AGAINST ACCESSED WAKE FEATURES)

This prediction technique is similar to prediction technique (6) but is based on
5 accessed features rather than visited features. It should be noted that since the
location-to-feature translation unit 58 does not know whether any feature it identifies
is an accessed feature, the prediction unit works on the basis of the most-recently
accessed feature as identified in the prediction assist data it receives from the user.

- 10 Rather than predicting the next feature to be accessed, the previous visit data can be
organized to enable the unit 58 to predict the next feature to be visited (and thus
feature items likely to be needed) or the next feature item visited / requested for
delivery (or delivered) / accessed.

15

9. SAME MOST-RECENTLY ACCESSED FEATURE HISTORY - PREDICTION OF FEATURE/
FEATURE ITEM

This prediction technique is similar to prediction technique (7) but is based on
accessed features rather than visited features. Again, rather than predicting the next
20 feature to be accessed, the previous visit data can be organized to enable the unit 58
to predict the next feature to be visited (and thus feature items likely to be needed) or
the next feature item visited / requested for delivery (or delivered) / accessed.

- 25 It will be appreciated that since each feature item can be represented by a respective
feature, where the foregoing prediction techniques involve features the same techniques
can be applied directly to feature items provided the latter have any required associated
parameter data, such as location. Thus, the prediction techniques (1), (2), (3), (4a) and (5a)
which all involve determining the closeness of features to a location or track, can equally
30 be implemented by determining the closeness of individual feature items to a location or
track. Similarly, techniques (6) to (9) can be applied by accessing the previous-users visit
histories in respect of the same visited/accessed feature item / feature-item tail rather than

the same visited/accessed feature / feature tail (in this context, a “visited” feature item is one where the user has visited the location associated with the location).

Where an above-described prediction technique is based on determining the probability of visiting / accessing a feature, then instead of using this probability as the probability of usage of all the feature items concerned, it is alternatively possible to set the usage probability of each feature time individually by weighting the feature-related probability according to the relative popularity (in terms of actual presentation to the user) of the item concerned with respect to other items associated with the same feature— provided, of course, that data about relative popularity is made available to the prediction unit 58.

All of the above prediction techniques can be implemented fully in service system 35, split in any appropriate manner between the service system 35 and the mobile devices 31, or fully in the mobile devices 31, even if based on the visit histories of previous users. Thus, for example, where prediction is done on the basis of previous visit histories but there is no service system 35, each mobile device can be arranged to store all its past visit histories and to supply them to other devices on request. As another example, the Figure 2 embodiment can be modified by arranging for the prediction unit 58 simply to provide the mobile device with the probabilities of features being visited / accessed, it then being up to the mobile device (in particular the cache manager 45) to translate features to feature items and request such items according to the probabilities associated with the corresponding features; this, of course requires the cache manager to have access to information about the association between the features and feature items and such information can conveniently be stored in memory 43. Rather than the cache manager 45 requesting individual items from the server 53 when effecting pre-emptive caching, it can supply a feature identifier to the server 53 which then returns all the feature items associated with the feature concerned.

Additional prediction techniques to those described above are also possible. Also, the above-described pre-emptive caching arrangement can equally be applied where the features items are being supplied to the cache 44 from a local storage device such as a DVD drive rather than from a remote resource over a wireless connection.

It is also possible to control loading of items into the cache 44 on the basis that they have not been identified as an item having a low probability of usage as determined using one of the above-described prediction techniques. In one implementation of this approach, the cache manager 45 is arranged by default to request from the server 53 all items associated with features within a predetermined distance of the user's current location (as determined, for example, by querying the feature data store 52); however, this default is overridden in respect of any item which, according to the prediction unit 52, has a probability of usage below a predetermined threshold value. In this example implementation, the prediction unit 52 is arranged to identify the low usage probability items based on the information received in the location reports 62 from the device 31, the identities of these items then being returned to the device 31 in a response message 65.

Other factors additional to item usage probability may be used to determine when an item should be loaded into cache. For example, the amount of free space in the cache can be used to control the threshold probability value below which items are not loaded into the cache – the fuller the cache, the higher this threshold is set.

Flushing the Feature Item Cache

An item retrieved by the mobile device 31 to the cache 44 will typically be retained in cache for as long as possible to be available for access by the user at any time including after the user has passed on from the feature with which the item is associated. However, since the size of the cache memory 44 will generally be much smaller than that required to store all available feature items, it will usually be necessary to repeatedly remove items from the cache during the course of a visit to make room for other features items. Items to be flushed from the cache are identified on the basis of a prediction-based indication of what items are unlikely to be needed again.

The above-described prediction techniques used for determining the probability of usage of feature items can also be used in determining whether a cached feature item should be flushed from the cache. The usage predictions can be used in any one or more of the following ways:

1). The item IDs and usage probabilities included in the second part 67 of the response message 65 returned to the mobile device 31 can be used to indicate whether an item has a probability of usage insufficiently high to justify it being maintained in the cache. This can be done by setting a probability threshold, known to the cache manager 45, below which items are to be flushed from the cache (or, optionally, they can be retained if there is no need to free up cache space). Of course, typically (though not necessarily) the second part of the response message 65 will only give the usage probabilities of a limited number of the available feature items, these being the items with the highest usage probabilities; in this case, an assumption is preferably made that all items not appearing in the second response part 67 have a usage probability below the aforesaid threshold and can therefore be flushed from the cache.

2). The prediction unit 58 can be arranged to include in a third part of the response message 65 the item IDs and usage probabilities of items which have low usage probabilities. The cache manager 45 uses these item IDs and usage probabilities to determine whether or not to flush the corresponding items from cache (if present). This determination is made, for example, by reference to a probability threshold below which items are to be (or can be) flushed. In fact, if the prediction unit 58 knows the value of this predetermined threshold, it can simply include in the third response part the IDs of items falling below this threshold. Of course, there may be a very large number of items with a low probability of usage – in particular, items associated with features that are distant from the user. Preferably, therefore, the prediction unit 58 restricts its determination of items with low usage probability by a filter adapted to tend to exclude items that are unlikely to be in the cache memory 44 of the user's mobile device. This filter can be based on the user's feature tail (only items associated with features in this tail have their usage probabilities assessed to ascertain if they are low), or the user's location tail (only items associated with features within a certain distance of this tail have their usage probabilities assessed to ascertain if they are low). Applying such a filter leaves open the possibility of a low usage probability item being retained in cache where that item is associated with a feature excluded by the filter. This can be avoided by arranging for the cache manager 45 to apply the same filter to all items in cache and then to flush all items

not passing the filter, except those items identified in the first part of the response message 66 (or, optionally, those items associated with the same features as the items identified in the first response part). For the cache manager to operate in this manner it will need to know which items are associated with which features and this association data can be downloaded to memory 43 at the start of the visit.

- 3) The cache manager 45 can be arranged to request the prediction unit to provide it with the probability of usage of individually identified items currently in the cache and then determine whether or not to flush each such item based on the usage probabilities returned by the prediction unit 58 (for example, if an item has a usage probability below a predetermined threshold value, the item is flushed). In making its request, the cache manager 45 sends the prediction unit 58 any data that the latter needs to apply the prediction technique being used ; this data corresponds to some or all of the data included in the most recent location report made by the mobile device 31. Depending on the prediction technique being used, the prediction unit 58 may need to extend the scope of its prediction operations to encompass the items concerned. Thus, for example, implementation of prediction technique (1) for a specified item requires first that the corresponding feature to be identified (in order to determine the location of the item) and then the probability of usage is determined - however, if the normal radius used in calculating probabilities according to this technique (30 meters in the example given above) is insufficient to cover all features, a greater radius - sufficient to cover all features - should be used for determining the usage probabilities of individual items. A preferred alternative is to assign a probability of zero to all items outside the normally used radius; in a similar manner, in techniques (2), (3), (4a) and (5a) the scope of enquiry can be limited to a next portion only of the planned / predicted / probable track of the user, any features not on or adjacent this next track portion being given a usage probability of zero. For prediction techniques that use count data in a manner analogous to that represented by table 140 in Figure 15 (such as techniques (6), (8) etc.), determining the usage probability of any particular item (or possibly, its associated feature) is simply a matter of accessing the correct table row to look up the count value for the item/feature concerned as the next item/feature and using this count and the total count for the row concerned to determine the usage probability. Adaptation of the

other prediction techniques for determining the usage probabilities of individual items will be apparent to persons skilled in the art.

5 A number of variants are possible to this third approach – for example, the cache manager 45 can provide feature identifiers to the prediction unit which then returns the probabilities of each of those features being visited / accessed; the cache manager then uses this information to make its determinations about whether to flush items associated with those features (the cache manager having access to information about the association between items and features). Another possible variant is to include in
10 each (or selected) location reports, the item IDs of items in cache, the prediction unit then returning the usage probabilities of these items in a third part of the response message 65. A further variant is to restrict the enquiry about usage probabilities of items in the cache 44 to those items that have not been recently accessed by the user (as indicated by a last-access timestamp associate with each item).

15 All of the foregoing cache-flushing arrangements can be implemented fully in the mobile devices 31 (even if based on the visit histories of previous users), split in any appropriate manner between the service system 35 and the mobile devices 31, or implemented fully in the service system 35 (apart from the actual operation of flushing identified items).

20 Other factors additional to item usage probability may be used to determine when an item should be flushed from the cache. For example, the amount of free space in the cache can be used to control the threshold probability value above which items are retained in the cache – the fuller the cache, the lower this threshold is set. Another factor that can be taken
25 into account is the time that has elapsed since an item of interest was accessed or, if no accesses have been made, since the item was first loaded into the cache; this factor can be used, for example as a weighting for a usage probability determined for the device – the longer the elapsed time, the smaller the weighting.

30 As regards when cache flushing is effected, this can be done (or tested for) each time a response message is received, or regularly as part of a garbage collection strategy, or simply when space is needed in the cache.

Transforming Cached Items

As an additional or alternative strategy to flushing items from cache, the cache manager can be arranged to increase the available space in the cache by transforming at least
 5 selected items already in cache 44 so that each takes up less space (fewer bytes of cache memory) than before. The selection of items to be transformed in this way is preferably done using the same prediction techniques and approaches to using the resulting usage probabilities as discussed above in relation to cache flushing; the timing as to when transformations are done can also be the same as discussed above for cache flushing

10 Furthermore, in selecting items for transformation, other factors besides usage probability can additionally (or, indeed, alternatively) be taken into account; possible other factors include cache free space and time since last access (or time since loading if no accesses have been made).

15 In a preferred arrangement, when the probability of usage of a cached item falls below a first level it is subject to transformation to take up less cache space; if the probability of usage of the item should fall below a second level, less than the first level, the item is flushed from cache (regardless of whether or not it has been previously transformed to take up less cache space.

20

The nature of the transformation to which an item is subject will generally be either compression using any of a variety of standard techniques, or a deliberate degradation where the available presentation quality of the item is traded for a reduction in the amount of cache space occupied by the item. Degradation of the item in step (b) can be effected,

25 for example, by at least one of:

- where the item comprises a sampled-media stream, reducing the sample rate and/or the number of bits used to represent each sample;
- selectively removing whole portions of the item (for example, replacing a long audio recording with just the first few seconds of its length);
- 30 - where the item is an image, reducing the resolution of the image;
- changing the format in which the item is represented.

The transformed item replaces the un-transformed version of the item in cache 44.

In a preferred embodiment, each cached item has an associated flag which is set when the item is transformed. This enables the cache manager 45 to tell whether an item has been transformed without examining the item in detail.

5

The reason to transform an item rather than flush it from cache is that it remains available to the user without the delay involved in having to fetch it from the server 53. For many applications, quick access to a reduced version of a media object item will be preferable to the user to slower access to the full version of the item.

10

In one preferred embodiment, upon an item being accessed for presentation to the user, it is retrieved from cache and presented without delay. At the same time as (or just before or just after) the item is first presented, the cache manager 45 checks the flag associated with the item or otherwise makes a determination as to whether the item has been transformed from the version originally received from the server 53. If the item has been transformed, the cache manager 45 requests the original version from the server again and when this version is received, it is substituted for the transformed version of the item being presented to the user. Where the item concerned is a streaming media item, the newly-received original version of the item is accessed for presentation at the point in the media stream currently reached during presentation of the transformed version of the item to the user.

20

Whilst it may normally be expected that an item loaded into cache 44 from server 53 will spend some time in cache in its un-transformed form before being selected for transformation, this is not necessarily the case. More particularly, an item received at the mobile device may already satisfy the condition set for selecting an item for transformation; in this case, the item is transformed immediately it is received.

25

Distributed cache

When the user of a mobile device 31 arrives at a new feature, it is likely that the mobile devices of other users already present at or near the feature will already have relevant feature items in their caches as a result of having accessed these items earlier. This

30

likelihood can be exploited by arranging for a mobile device 31 that wants to load a particular feature item (referred to below as the 'requesting' device), to try first to obtain it from other mobile devices physically nearby and only if this is not possible, resort to accessing it from the item server 53 (or from whatever its original, non-device, source may be). This has benefits for the item server 53 (or other source) in reducing its load; furthermore, if, as is preferred, a separate communication mechanism is used for device-to-device communication (such as Bluetooth short-range radio links) as compared with device-server communication (the WLAN in the embodiment of Figures 1 and 2), then the bandwidth loading on the network used by the server 53 to distribute items is also reduced. The response time for a requesting device to receive a requested item will typically also be reduced.

In effect, the caches 44 of nearby mobile devices serve as a distributed cache of feature items for the requesting device.

15

This way for a requesting device to obtain a needed item can be applied both in the case where pre-emptive caching is not done by the device (an item only being requested when needed for presentation to the user), and in the case where pre-emptive caching of items (as, for example, described above) is being effected. In this latter case, the degree of pre-emption (that is, how many items are loaded pre-emptively) is preferably reduced as compared to where nearby devices do not provide an available distributed cache of feature items.

20

It will be appreciated that whilst it is not essential for the requesting device to make its own cache reciprocally accessible to other devices in order for it to benefit from the distributed cache of feature items offered by other devices, this will normally be the case. It will also be appreciated that even where a device only requests an item when needed for presentation to a user, it will still normally be provided with a cache to enable it to temporarily retain the accessed item in case the user wishes to refer back to it later on.

25
30

The requesting device seeks to retrieve a needed feature item first from devices that are nearby because the feature items have location relevancy and so it is the nearby devices

that are most likely to have the item of interest in cache. How the request for an item is limited to nearby devices can be done in a number of ways. For example, this can be done implicitly by the requesting device using a short-range communication technology to send out its request so that only nearby devices receive the request. Alternatively, where the locations of the mobile devices are known as is the case for the mobile devices of the embodiment of Figures 1 and 2, then a determination can be made (for example, by the service system 35) as to which other mobile devices are close to the requesting device; these devices can then be contacted to ascertain if any of them have the needed item. Another way of delimiting the nearby devices is by specifying in advance a group of devices that are likely to be near each other (such as a tour party) – in this case, the devices specified are those likely to be nearby rather than those that are actually nearby.

Rather than determining the target group of devices to be contacted for the item of interest on the basis of closeness to the requesting device, this target group can be determined on the basis of which devices are close to the location associated with the item of interest.

In fact, rather than limiting the enquiry for the item of interest to the group of devices that are near (or likely to be near) the requesting device or the location associated with the item of interest, it is alternatively possible to arrange for the item server 53 or other functionality to track for each item which device or devices currently hold that item in cache, or have received and are likely still to have the item in cache. In this case, the requesting device can be arranged to contact mobile devices likely to have the item of interest regardless of whether these mobile devices are nearby (though this constriction may, of course, be implicitly applied where device-to-device communication is by short-range communication means).

Whether the requesting device seeks to retrieve a needed item from a device that is likely to be nearby or from a device that is likely to have the item for some other reason, additional filter(s) can be applied as to the devices used to supply the item. Thus, for example, the requesting device may be arranged to contact all nearby devices using a short-range communications means but only accept to receive the item from a device that is also

part of the same tour party (as indicated by an indicator stored as part of the visit data in memory 43).

Two example implementations of a requesting device arranged to seek to retrieve a needed item first from nearby devices will now be described with reference to the simplified process flow charts Figures 16 and 17 respectively. In both implementations, it will be assumed that the mobile devices have short-range communications means, such as a Bluetooth radio system, by which they can communicate with each other, these short range communication means being additional to the communication means used to communicate with the server system 35.

In the first implementation (Figure 16), the requesting mobile device 31 on determining that it needs a particular feature item, starts off (block 150) by using the short-range communication means to contact all nearby devices (that is, devices within communication range) to ask whether any of these devices have the item of interest in cache (block 151). Upon a contacted mobile device receiving this item availability request (block 152), it checks its cache contents and sends a response (block 153) indicative of whether or not it has the item of interest. The requesting device collates the responses (if any) it receives (block 154) and determines if there are any positive responses (block 155). Assuming at least one positive response is received, the requesting device now requests (block 156) the positively-responding device, or a selected one of the devices if more than one device has responded positively, to send the item of interest. The positively-responding device concerned, on receiving the item request (block 157), returns the item to the requesting device (block 158). The requesting device receives and stores the item (block 159). At this point the process of fetching the item is complete (block 160).

If the requesting device determines at block 155 that it has failed to receive any positive response to its item availability request within a predetermined timeout period, it contacts the item server 53 (block 161) for the item of interest. On receiving this request (block 162), the server 53 returns the item (block 163) to the requesting device which receives and stores it (block 159).

As a variant of this first implementation, the requesting device can be arranged to send its availability request to specific devices rather than as a broadcast message to all devices within range. To this end, the requesting device keeps a list of devices it can contact. This list can be a fixed list set, for example, at the start of the visit and comprising mobile devices of the same tour party. Alternatively, the list can comprise mobile devices known to be close to the requesting device as a result of a comparison of their location with that of the requesting device; such a list has, for example, been compiled by the service system based on the device locations it has received in the location reports from the mobile devices, an updated version of this list being periodically sent to the requesting mobile device (for example, as part of a response message 62). However the list of devices to contact is compiled, the requesting device may contact each device on the list in turn until a positive response is received or may send a multicast message to all devices on the list where the communications means concerned supports multicasting.

15 In the second implementation (Figure 17), the requesting mobile device 31 on determining that it needs a particular item, starts (block 170) by contacting the item server (block 171), or some other service-system functionality, to find out which other mobile device it should contact to obtain a copy of the item. The item server on receiving this request (block 172), determines which device if any has, or is likely to have the item of interest (block 173) and returns the identity of this device (or devices) to the requesting device (block 174). If the item server determines at block 173 that no device has, or is likely to have, the item, then the server itself sends the item to the requesting device (block 175) which receives and stores the item (block 176) thereby completing the item retrieval process (block 177).

25 The item server can make its determination in block 173 about which devices have, or are likely to have, the item of interest in any of a number of different ways. Thus, as already indicated, the server can use a list of devices known to be related to the requesting device as a result of them all being associated with the same visit party. Alternatively, the server can query functionality of the service system that knows which other devices are close to the requesting device (the pheromone trail subsystem can provide this functionality where it stores the virtual markers left by devices in non-aggregated form against device identity). In a further alternative, the item server or associated functionality is arranged to contact all

other mobile devices, or just those near the requesting mobile device, with an item availability request, the positively-responding devices then being identified back to the requesting device in block 174. A yet further alternative involves the item server or associated functionality keeping track of which devices have the item of interest in cache
 5 or have previously received the item and are likely still to have it in cache.

Where the item server determines that at least one mobile device has, or is likely to have, the item of interest to the requesting device, the identity of the or each of these devices is returned to the requesting device which receives and temporarily stores these device
 10 identities (block 178). The requesting device then determines what action to take next (block 179) which in this case is to contact a first one of the identified devices, using the short-range communication means, to ask it for the item of interest (block 180). The contacted device receives the request (block 183) and decides what action to take (block 184) in dependence on whether or not it has the item of interest in cache. If the contact
 15 device has the item of interest, it now returns it to the requesting device (block 186) which receives and stores it (block 176) thereby ending the retrieval process. However, if the contacted device does not have the item of interest, it returns a negative response to the requesting device (block 185). Upon the requesting device receiving this negative response (block 178), it determines what further action to take. If there remains a device not yet
 20 contacted in the set of devices identified by the server, the requesting device now contacts the next device of this set. However, if all of the identified devices have been contacted and all have responded negatively (or not at all within a timeout period), then at block 179 the requesting device determines that it must ask for the item from the item server and this it now does (block 181). The item server, on receiving such a request (block 182), responds
 25 by sending back the item (block 175).

In both the first and second implementations (Figures 16 and 17), it is possible to arrange for a contacted mobile device that does not have the item of interest, to participate more actively in the process of finding a device with the item of interest. More particularly,
 30 where a contacted device has an indication as to what other device may have that item, it can be arranged either to pass the identity of that device back to the requesting device or to pass on the request from the latter to that other device. A contacted device can have an

- indication about what other device may have the item of interest in a variety of ways; for example, the contacted device may have itself passed the item of interest on to another device, or the contacted device may have received the item from another device (the contacted device having subsequently flushed the item from its cache). Another possibility, in the case of the Figure 17 implementation, is that the contacted mobile device received a list of devices that had, or may have had, the item of interest. In all these cases, provided the contacted device has kept a record of the appropriate information, it can assist the requesting device in finding another device holding the item.
- As already mentioned, it is possible to arrange for the item server or other functionality to keep track of which mobile devices have, or are likely to have, each feature item. A simple way of doing this is for the server (or other functionality) to keep a record of the devices to which the item server has served each item in an immediately preceding time window (for example, of five minutes). This record is preferably supplemented by information about the device-to-device transfers of items – for example, either the sending or receiving device can be required to inform the item server (or other functionality) about such a transfer; including the time it was done. The record can be further supplemented by having the devices inform the item server (or other functionality) whenever they delete an item from cache and if this is done, the time window requirement can be removed (or set with a much longer duration). In this way a fairly comprehensive record can be kept about which devices are holding which items. Another approach would be to have each mobile device regularly report what feature items are present in its cache; this information can be incrementally updated, between the regular full reports, as items are added and flushed.
- Arranging for the item server (or other functionality) to keep track of which mobile devices hold which items, not only permits a requesting device to be pointed quickly at a device with the item of interest, but also enables the item server to pre-emptively push copies of feature items relevant to a particular zone (for example, a room of the hall 10) to devices in that zone. Thus the item server can determine what items relevant to a zone are not held by any of the mobile devices currently in that room and then take action to have these items copied to devices in the zone. This can be done directly by giving the item server the capability to push items to devices or, where pre-emptive caching is implemented by the

devices on the basis of item usage probabilities, indirectly by artificially raising the probabilities of the items not yet cached by devices in the zone – for example, the usage probabilities of these items can be set to unity to cause them to be downloaded immediately. Where this latter approach is used then, in order to distribute such items

5 between the devices in the zone, for each device only certain of the items not cached in devices in the zone would have their usage probabilities raised in this way. The direct or indirect pushing of items relevant to a zone to devices in that zone can be applied to a subset of the items relevant to the zone rather than to all relevant items; in particular, infrequently used or especially large items can be omitted.

10

In fact, it is possible to arrange for the item server to try to put all items relevant to a zone into the caches of devices in the zone without the need to track what devices hold what items. This can be done by arranging for every device in the zone to take a proportion size of the relevant items, this proportion varying according to the number of devices currently

15 in the zone; whenever a mobile device enters the zone it is pushed its proportion of items, the actual items concerned being identified by taking the items in sequence and continuously cycling through all the items relevant to the zone. The proportion of items allocated to each device is preferably judged in terms of the amount of memory space taken up by the items rather than simply on the basis of the number of items.

20

A similar, but more restricted, effect can be provided by requiring that a mobile device moving from one zone to another pass on copies of the items relevant to the zone it is leaving to devices in the zone so that at least one copy of each such item is held in the cache of a device in the zone.

25

Where a track is kept of which mobile devices hold which items, rather than this task being carried out by the item server or other functionality at the server system, tracking can be effected for each zone by one of the mobile devices in the zone, the devices in the zone reporting to that device when they receive and flush items. The mobile device allocated

30 this task can then not only serve to identify to requesting devices which nearby device holds a particular item, but can also be arranged to cause items not locally cached to be pushed directly or indirectly to devices in the zone concerned. When the device allocated

this task passes to another zone, the responsibility for carrying out the task is passed to another device still in the zone, either by the exiting device itself or by functionality of the service system.

- 5 As already indicated, whilst it is preferred that the device-to-device transfer of feature items is effected by a separate communications mechanism to that used for server-to-device item transfers, this is not essential. One situation where advantages are still to be gained by having devices trying first to obtain an item of interest from a nearby device rather from the server even though the device uses the same communication mechanism for
- 10 communicating with the server and other devices, is where several wireless LANs are being used to cover different parts of a space both for server-device communication and for device-device communication. In this case, even though a requesting device may take up bandwidth on one wireless LAN whilst receiving an item from a nearby device, the server can transfer a different item to another device communicating with it over a different one
- 15 of the wireless LANs.

Coordinated Consumption

- Often a user visiting the exhibition hall 10 will be doing so as a member of a party, be it a family group, a tour party or some other group. Where at least some members of the party
- 20 have respective mobile devices, the situation is likely to arise that one member with a mobile device accesses a feature item that they then wish to share with the other party members with mobile devices; indeed, this may the case not only for feature items but for any data item available from the service system 35 or other source accessible via the
- 25 communications infrastructure exemplified in the embodiment of Figures 1 and 2 by wireless LAN 36. Where the data item is a simple media object such as image, then this is can be achieved by the accessing member passing on the identity of the item to the other members either verbally or possibly by a message sent from their mobile device to the other mobile devices associated with the party.

30

However, where the item concerned is a streamed media object (in particular, audio and/or video streams), simply having each mobile device independently access the object will

result in uncoordinated consumption of the object at each device.

An arrangement for providing for coordinated consumption of the streamed media object is illustrated in Figure 18. In this Figure, the streamed media object 200 held on server 53 is depicted as being streamed (arrow 201) to the mobile device 31A of a first user 30A. Upon the user 30A wished to share the experience provided by this media object with another user 30B, the user 30A causes their mobile device 31A (for example, by pressing a dedicated button of user interface 48) to send a "share" message 202 to mobile device 31B (see arrow 203). This message is, for example, sent via the wireless LAN 36 using the addresses of the other mobile devices of members of the same party, these addresses having been previously stored in the visit data memory 43 of device 31A; alternatively, the device 31A may simply use a short-range communications means, such as a Bluetooth radio system, to send the message 202 to any device that is nearby (this approach, though less secure, is generally acceptable in an environment such as the exhibition hall because the media object itself is not confidential). The message 202 includes both an identifier of the media object 200 currently being accessed by the mobile device 31A (for example, in the form of an item URI – Uniform Resource Indicator), and an indicator of the current position reached by the user 30A in consuming the streamed media object 200 (for example, frame number, time point, etc.).

The mobile device 31B on receiving the message 200 records its time of receipt and then, with or without specific consent from the user 30B, contacts the server 53 (arrow 204) and accesses the media object 200, the latter being delivered to device 31B as a media stream (arrow 205) independent of the media stream (arrow 201) being consumed by the device 31A. Since streamed media is generally delivered at a rate greater than its rate of consumption with the receiving device buffering the received but not yet consumed portions of the stream, the receipt by device 31B of the media object 200 can catch up with consumption by the first device 31A. Thus, mobile device 31B is arranged to delay rendering (presenting) the streamed media object until delivery of the object 200 reaches the position in the stream indicated in the message 200 plus any advance that the device 31B determines it should add to allow for further consumption of the media object by the first device 31B between when the position indicator was inserted in the message 202 and

the start of rendering by the second device 31B. This advance, where applied, can comprise one or more of the following components:

- a component taking account of the message assembly and transmission time from the mobile device 31A until received at the mobile device 31B (this can be a preset approximation);
- a component taking account of the time between message receipt at device 31B and the start of receipt at the device of the media object 200 (as timed internally by the device 31B);
- a component taking account of the time to move through the media object being streamed to the device 31B to reach a position corresponding to the then current point of consumption at device 31A (this will generally involve a predictive process).

Rather than the device 31B only starting rendering (presenting) the media stream to user 30B when the estimated position of coordination with device 31A is reached, the device 31B can be arranged to render the preceding portion of the stream in fast time – that is, presentation to the user will appear to be at high speed though typically in practice this is done by rendering only spaced samples of the stream (for example, every n th frame of a video sequence). Normal rendering of the stream is then started when the consumption of the media stream at device 31B reaches the estimated position of coordination with device 31A.

20

Where the server 53 is operative to initiate streaming of the media object 200 at a specified offset into the object media stream, then the mobile device 31B is preferably arranged to take advantage of this capability by requesting streaming to start at the nearest available (lagging) offset to a point that the device 31B predicts will be the current point of consumption by the first device 31A by the time the second device 31B is ready to render the media stream from the server 53. A similar effect can be obtained by the second device 31B informing the server of the identity of the first device 31A and then arranging for the server to commence delivery of the media object to the second device 31B from a position corresponding to the current delivery position to the first device 31A less an amount corresponding to the cache size allowed in the first device 31A for caching the media object as it is streamed. This latter approach is primarily of interest if this cache size is relatively small as compared to the overall run time of the media object.

30

Of course, the first user 30A may decide to pause consumption of the media object 200 at the time the user decides to share experiencing the object with the other members of the same party, this pause being to enable the other users to access the media object at the earliest possible coordinated position for going forward from the first user's current position. If the first user pauses the streaming of object 200 at the time of sending message 202, this is preferably indicated in the message 202 and enables the receiving device 31B to dispense with adding any advance to the position indicated in the message. When the second device 31B becomes ready to render the media object from that indicated position, it signals to the first device 31A to resume consumption of object 200 and then itself starts consumption from the indicated position (possibly with a small delay corresponding to the time needed for the first device to react). In fact, rather than the second device 31B starting consumption based on when it sends a signal to the first device indicating that it is ready to do so, the second device 31B can be arranged to wait for a "go" signal back from the first device before starting. This enables the first device 31A to ensure that all devices which are to participate in coordinated consumption, are ready to do so before consumption begins (in other words, the device 31A waits to receive signals from all expected devices indicating they are ready to start consumption before it sends out the "go" signal).

It is possible to arrange for the second mobile device(s) 31B to assume that the first device 31A will have paused the media stream at the time of sending the message 202 (indeed, such a pause can be enforced at device 31A in coordination with sending of the message 202); in this case, the message would not need to include a "pause" indication. If this default arrangement is used but continued consumption at the instigating device 31A is still permitted, then in cases where device 31A continue its consumption, this is preferably indicated in message 202 to permit the device 31B to add an appropriate advance as already discussed.

One or both of the devices 31A and B can be arranged to send further coordination signals at any time during the course of consumption of the media object 200 by the devices, in order to compensate for drift in consumption rates. Thus, for example, the device 31A as the instigating device, can be arranged to send periodic coordination signals to the device

31B which indicate the current position reached by the device 31A; the device 31B uses this information either to jump backwards/forwards in the stream it is receiving or to make appropriate adjustments to its rate of consumption of the media stream to return to synchronization with consumption by device 31A (for example, by the time the next
 5 coordination signal is expected to be received).

As well as the sending of periodic coordination signals, one or both of the devices can be arranged to send change signals in correspondence to its own changes in position in the media stream (other than resulting from normal progression therethrough) and/or changes
 10 in its own progression through the media stream (such as, without limitation, rewind and fast forward), these change signals including an indication of the new position/progression mode to be adopted by the receiving device.

Rather than (or additional to) the current position reached in the media object 200 by
 15 device 31A being included in message 202, this information can be passed to the device 31B in a separate message sent in response to device 31B indicating that it is starting to receive the media object stream from server 53. This approach enables a more up-to-date current position to be used by the device 31B in determining when to start normal rendering of the media object 200.

20 Whilst coordinated consumption could be effected in the above-described manner for static devices separate by large distances, it is expected that the above-described method is most suitable where at least one of the devices is a mobile one and the devices have been brought into close proximity.

25

Variants

It will be appreciated that many variants are possible to the above described embodiments of the invention. For example, although in all the embodiments described above, all feature items have originated from the same source, namely, item server 53, it is also possible to
 30 provide for multiple item sources each holding a respective subset of the items. In this case, the item identifier associated with each item can be arranged to indicate directly the source from which the item can be obtained, or some other mechanism can be employed to

direct an item request to the appropriate source. The multiple item sources effectively form a distributed item server.

5 As already noted, the distribution of functionality between mobile devices and the service system is not limited to the distributions described above since the availability of communication resources makes it possible to place functionality where most suitable from technical and commercial considerations. Furthermore, in the foregoing reference to a mobile device is not to be construed as requiring functionality housed in a single unit and the functionality associated with a mobile device can be provided by a local aggregation of
10 units.

The above described methods and arrangements are not limited to use in exhibition halls or similar public or private buildings; the methods and arrangements disclosed can be applied not only to internal spaces but also to external spaces or combinations of the two.
15

66
CLAIMS

1. A method of providing information about a real-world space, comprising the steps of:
 - (a) as each of at least one first user moves through said space, a succession of virtual markers is deposited by a mobile device carried by the user in a digital representation of the space where they are stored to indicate locations visited by the user in the space;
 - (b) on a second user moving through the space, the virtual markers deposited by the said at least one first user are accessed to provide to a mobile device of the second user information for facilitating use of the space by the second user.
2. A method according to claim 1, wherein there are multiple first users and their associated virtual markers are aggregated, in dependence on their associated locations, either in step (a) at the time they are recorded in said digital representation, or in step (b) at the time they are used to provide said information for facilitating use of the space by the second user.
3. A method according to claim 2, wherein the virtual markers associated with said first users are of multiple types, the aggregation of markers being done by type.
4. A method according to any one of the preceding claims, wherein the virtual markers have an initial strength value and the strength value associated with a deposited marker, either individually or in aggregation with other markers, is caused to decay with time.
5. A method according to any one of the preceding claims, wherein said virtual markers associated with a first user are deposited automatically by the mobile device of the user at one of:
 - predetermined intervals of time;
 - predetermined intervals of distance; or
 - predetermined locations in said space.
6. A method according to claim 2 or 3, or to claim 4 or 5 when dependent on claim 2 or 3, wherein step (b) involves using the virtual markers associated with said first users to

determine a path through the space by following or avoiding peaks or troughs in the virtual landscape formed by the aggregations of markers.

7. A method according to claim 1, wherein the virtual markers deposited by the mobile device of a said first user have associated data indicative of the user concerned, step (b) involving providing information about the path taken by that first user by identifying the virtual markers associated with that user.

8. A method according to claim 7, wherein the virtual markers that have said associated data indicative of the user concerned, have an initial strength value and the strength value of each such marker is caused to decay with time; step (b) including using the relative strength values of the markers to determine the direction of progression of the user concerned along said path.

9. A method according to claim 2 or 3; or to claim 4 or 5 when dependent on claim 2 or 3, wherein step (b) involves using the virtual markers associated with said first users to predict a next location for the second user having regard to that user's current location, this predicted location then being used to provide to said mobile device, as said information, either the identify of media items associated with that predicted location or the items themselves.

10. A method according to claim 2 or 3, or to claim 4 or 5 when dependent on claim 2 or 3, wherein in step (a) a said virtual marker is deposited whenever a said first user reaches a location corresponding to an item of interest, step (b) involving using the virtual markers associated with said first users to provide information about the popularity of items of interest in said space.

11. A method of guiding a user along a target path, comprising the steps of:
(a) determining the position of the user relative to the target path; and

(b) providing respective audio cues to the user via left and right audio channels, these cues being indicative of the relative position determined in step (a) and varying in a complementary manner over at least a range of values of said relative position.

5 12. A method according to claim 11, wherein the same audio characteristic of the audio cues delivered by the left and right channels is varied with changes in said relative position over said range of values, this characteristic being varied in one sense for the cues delivered by the left channel and in the opposite sense for cues delivered by the right channel as said relative position changes.

10

13. A method according to claim 12, wherein said audio characteristic is at least one of frequency and volume.

15 14. A method according to claim 11, wherein the same audio characteristic of the audio cues delivered by the left and right channels is varied with changes in said relative position over said range of values, said audio characteristic being increased/decreased in magnitude for one said channel as the user moves away from the target path to the side of the path corresponding to that channel, and the said audio characteristic of the other channel being correspondingly decreased/increased in magnitude.

20

15. A method according to any one of claims 11 to 14, wherein said relative position is determined in step (a) in terms of the perpendicular distance between a centreline of the target path and the user's current position.

25 16. A method according to any one of claims 11 to 15, wherein the audio cues are only provided whilst the user is within a predetermined distance of the target path centreline.

17. A method according to any one of claims 11 to 16, wherein a characteristic of the audio cues is varied in dependence on distance moved along the target path.

30

18. A method according to any one of claims 11 to 17, wherein the method includes the further steps, carried out when the user is moving, of:

- determining the angle between the user's direction of moving and the direction of pointing of the target-path centreline;
- where this angle is greater than a predetermined magnitude, providing an audible indication to the user of the direction of pointing of the target-path centreline.

5

19. A method according to any one of claims 11 to 18, wherein the method includes the further steps, carried out when the user is stationary, of:

- determining the angle between the user's direction of facing and the direction of pointing of the target-path centreline;
- 10 - where this angle is greater than a threshold magnitude, providing an audible indication to the user of the direction of pointing of the target-path centreline.

20. A method according to claim 18 or claim 19, wherein said audible indication is one of:

- a synthesized sound source rendered to appear to emanate from a location along the direction of pointing of the target-path centreline;
- 15 - a sound signal, independent of said audio cues, provided in the sound channel corresponding to the side of the target path to which the direction of moving or direction of facing, as the case may be, is pointing; and
- a variation in said audio cues indicative of the side of the target path to which the direction of moving or direction of facing, as the case may be, is pointing.
- 20

21. A method of guiding a user along a target path, comprising the steps of:

- (a) determining the position of the user relative to the target path; and
- 25 (b) rendering an audio beacon through audio output devices carried by the user such that the beacon appears to the user to emanate from a location in a direction at least approximating the direction of the target path onward from the user's current position.

22. A method according to claim 21, wherein the said location from which the audio beacon appears to emanate is changed each time the user approaches or arrives at the location whereby the audio beacon appears to move stepwise ahead of the user as the user progresses along the target path.

30

23. A method according to claim 22, wherein each successive location of said audio beacon is determined by determining a segment onward from the user's current position of a piecewise linear approximation to said target path, and setting said location at or relative to the end of this segment.
24. A method according to claim 21, wherein step (b) involves effecting at least a partial piecewise linear approximation of the target path and determining said location from which the audio beacon appear to emanate at or relative to the end of a segment of that approximation on or closest to which the user is currently positioned.
25. A method according to claim 21, wherein steps (a) and (b) are repeatedly continually with the location associated with the audio beacon being changed to appear to be at a substantially constant distance ahead of the user as they move along the target path at least over a substantial portion of the latter.
26. A method according to claim 21, wherein in step (b) one or more further audio beacons are each rendered through said audio devices to appear to the user to emanate at a respective further location selected such that the audio beacons together form a succession of beacons with each beacon being successively further down said target path and the or each said further beacon serving to indicate a direction to be followed from a preceding said beacon.
27. A method according to claim 26, wherein step (b) involves effecting at least a partial piecewise linear approximation of the target path and determining the locations from which the audio beacons appear to emanate at or relative to the end of respective successive segments of said approximation.
28. A method according to claim 26 or 27, wherein as the user approaches or arrives at the first audio beacon in said succession that beacon is removed, a new further beacon being added to the end of succession in time proximity to the removal of the first beacon in said

succession, this removal and addition of audio beacons being repeated as the user moves along the target path.

29. A method according to any one of claims 26 to 28, wherein an audible characteristic of said audio beacons is varied between beacons to indicate the order in which they occur along said path.

30. A method according to claim 29, wherein the audio beacons sound in the order they occur in said succession and in a cyclic manner.

10

31. A method of managing a cache of a mobile device carried by a user, the cache being used for storing items associated with locations in a real-world space being visited by the user; the method comprising the steps of:

- 15 (a) determining the probability of usage of an item in dependence on the user's progress around the space;
- (b) changing the contents of the cache by adding or removing an item on the basis of the determination carried out in step (a) in respect of that item or other items.

20 32. A method according to claim 31, wherein in step (a) said probability of usage is determined on the basis of the distance between the location associated with said item and the user's current location in said space.

25 33. A method according to claim 32, wherein step (a) includes reducing said probability of usage where said item is associated with a location lying in a wake region extending behind the user with respect to the user's progression through the space.

30 34. A method according to claim 31, wherein in step (a) said probability of usage is determined on the basis of the distance between the location associated with said item and the onward track from the user's current location of a planned route being followed by the user.

35. A method according to claim 31, wherein in step (a) said probability of usage is determined on the basis of the distance between the location associated with said item and the onward track from the user's current location as predicted on the basis of the user's recent movement in said space.

5

36. A method according to claim 31, wherein in step (a) said probability of usage is determined using visit history data of previous users that have visited the space, step (a) including identifying relevant visit history data for use in determining said probability of usage by matching the value of an indicator of the current user's progress around the space

10 with values of that indicator in said visit history data.

37. A method according to claim 36, wherein said indicator is the user's current location, the item usage probability determined in step (a) being determined by reference to the visit history data of previous users who have been in the same location as the user's current location, this determination using portions of said visit history data relevant to the track taken by previous users from the user's current location in order to determine an onward track for the user, the probability of usage of said item being derived on the basis of the distance between the location associated with the item and said onward track.

20

38. A method according to claim 36, wherein said indicator is the user's current location, the item usage probability determined in step (a) being determined by reference to the visit history data of previous users who have been in the same location as the user's current location, this determination using portions of said visit history data relevant to item usage

25 onward from the user's current location.

39. A method according to claim 38, wherein step (a) includes reducing said probability of usage where said item is associated with a location lying in a wake region extending behind the user with respect to the user's progression through the space.

30

40. A method according to claim 36, wherein said indicator is the user's recent movement in said space, the item usage probability determined in step (a) being determined by

reference to the visit history data of previous users whose movement to the user's current location corresponds to that of the user's recent movement, this determination using portions of said visit history data relevant to the track taken by previous users from the user's current location in order to determine an onward track for the user, the probability of
5 usage of said item being derived on the basis of the distance between the location associated with the item and said onward track.

41. A method according to claim 36, wherein said indicator is the user's recent movement in said space, the item usage probability determined in step (a) being determined by
10 reference to the visit history data of previous users whose movement to the user's current location corresponds to that of the user's recent movement, this determination using portions of said visit history data relevant to item usage onward from the user's current location.

15 42. A method according to claim 36, wherein said indicator is the identity of the item whose associated location has been most-recently visited by the user, the item usage probability determined in step (a) being determined by reference to the visit history data of previous users who visited the same item-associated location as the user's most-recently visited item-associated location, this determination using portions of said visit history data
20 relevant to item usage onward from the user's most-recently visited item-associated location.

43. A method according to claim 36, wherein said items are associated with virtual features each of which has an associated location in said space, the or each item associated
25 with a said feature having as its own associated location the location associated with that feature; said indicator of the current user's progress around the space being the feature most recently visited by the user, and the item usage probability determined in step (a) being determined by reference to the visit history data of previous users who visited the same feature as the user's most-recently visited feature, this determination using portions
30 of said visit history data relevant to item usage onward from the user's most-recently visited feature.

44. A method according to claim 36, wherein said indicator of the current user's progress around the space is the sequence of at least two items whose item-associated locations have been most recently visited by the user, the item usage probability determined in step (a) being determined by reference to the visit history data of previous users having a same sequence of visited item-associated locations as the sequence of item-associated locations most recently visited by the user, this determination using portions of said visit history data relevant to item usage onward from the user's most-recently visited item-associated location.
45. A method according to claim 36, wherein said items are associated with virtual features each of which has an associated location in said space, the or each item associated with a said feature having as its own associated location the location associated with that feature; said indicator of the current user's progress around the space being the sequence of at least the two features most recently visited by the user, and the item usage probability determined in step (a) being determined by reference to the visit history data of previous users having a same sequence of visited features as the sequence of features most recently visited by the user, this determination using portions of said visit history data relevant to item usage onward from the user's most-recently visited feature.
46. A method according to claim 36, wherein said indicator is the identity of the item most-recently accessed for presentation by the user, the item usage probability determined in step (a) being determined by reference to the visit history data of previous users who accessed for presentation the same item as most-recently accessed for presentation by the user, this determination using portions of said visit history data relevant to item usage onward from the user's most-recently accessed item.
47. A method according to claim 36, wherein said items are associated with virtual features each of which has an associated location in said space, the or each item associated with a said feature having as its own associated location the location associated with that feature; said indicator of the current user's progress around the space being the feature associated with the item most recently accessed for presentation by the user, and the item usage probability determined in step (a) being determined by reference to the visit history

data of previous users who accessed for presentation an item associated with the same feature as the item most-recently accessed for presentation by the user, this determination using portions of said visit history data relevant to item usage onward from the feature associated with the item most-recently accessed for presentation by the user.

5

48. A method according to claim 36, wherein said indicator of the current user's progress around the space is the sequence of at least the two items most recently accessed for presentation by the user, the item usage probability determined in step (a) being determined by reference to the visit history data of previous users having a same sequence of items
10 accessed for presentation as the sequence of items most recently accessed for presentation by the user, this determination using portions of said visit history data relevant to item usage onward from the user's most-recently accessed item.

49. A method according to claim 36, wherein said items are associated with virtual
15 features each of which has an associated location in said space, the or each item associated with a said feature having as its own associated location the location associated with that feature; said indicator of the current user's progress around the space being the sequence of at least the two features associated with items most recently accessed for presentation by the user, and the item usage probability determined in step (a) being determined by
20 reference to the visit history data of previous users having a same sequence of features with items accessed for presentation as the sequence of such features for items most recently accessed for presentation by the user, this determination using portions of said visit history data relevant to item usage onward from the user's most-recently visited feature.

25 50. A method according to any one of claims 38, 39, and 41 to 49, wherein said visit history data relevant to item usage is data about one of:

- the item or group of associated items next visited;
- the item next accessed for presentation or the group of items with which that item is associated;
- 30 - the item next delivered or requested for delivery to the cache.

51. A method according to any one of claims 30 to 50, wherein in step (b) an item is loaded into the cache, this item being one of:

- the item in respect of which step (a) is effected, the probability of usage of the item being determined as being above a threshold value for loading items in the cache; or
- 5 - an item not identified in a set of items have probabilities of usage, as determined by step (a), below a threshold value for loading items in the cache.

52. A method according to any one of claims 30 to 50, wherein in step (b) an item is removed from the cache, this item being one of:

- 10 - the item in respect of which step (a) is effected, the probability of usage of the item being determined as being below a threshold value for retaining items in the cache; or
- an item not identified in a set of items have probabilities of usage, as determined by step (a), above a threshold value for loading or retaining items in the cache.

15 53. A method of managing a cache of a mobile device carried by a user, the cache being used for storing items associated with locations in a real-world space being visited by the user; the method comprising the steps of:

(a) receiving an item at the mobile device, and

(b) degrading the received item to reduce the amount of cache space needed to store it, and

20 storing the degraded item in the cache instead of the un-degraded item.

54. A method according to claim 53, wherein as part of step (a) the received item is initially stored in the cache in an un-degraded form, the degrading of the item in step (b) being subsequently effected upon a predetermined condition concerning the item and/or the

25 mobile device becoming satisfied.

55. A method according to claim 54, wherein said predetermined condition is at least partially based on the probability of usage of the item as assessed from a determination of the probability of usage of that or other items having regard to the progress of the user

30 around said space.

56. A method according to claim 54 or claim 55, wherein said predetermined condition is at least partially based on the time elapsed since the item was last accessed by the user or, if not yet accessed, since the item was initially loaded into the cache.
- 5 57. A method according to any one of claims 54 to 56, wherein said predetermined condition is at least partially based on the amount of available cache space remaining.
58. A method according to claim 53, wherein step (b) is effected upon the receipt of the item at the mobile device in the event that a predetermined condition concerning the item
10 and/or the mobile device is already satisfied.
59. A method according to any one of claims 53 to 58, wherein the degrading of the item in step (b) is effected by at least one of:
- where the item comprises a sampled media stream, reducing the sample rate and/or
15 the number of bits used to represent each sample;
 - selectively removing whole portions of the item;
 - where the item is an image, reducing the resolution of the image;
 - changing the format in which the item is represented.
- 20 60. A method according to any one of claims 53 to 59, wherein the degraded item is stored in cache along with an associated flag indicative of its degraded form.
61. A method of presenting an item to a user of a mobile device where said item is stored in a cache of the mobile device according to the cache management method of any one of
25 claims 53 to 60, the item presentation method comprising the steps of:
- (a) upon the item being required for presentation to the user, retrieving the item from cache and presenting it to the user;
 - (b) determining whether the item is in a degraded form, this step being carried out before, at the same time as, or after step (a);
 - 30 (c) where step (b) indicates that the item is in a degraded form, requesting an un-degraded version of the item from an off-device resource and when received, substituting this version of the item for the degraded version being presented to the user.

62. A method of managing a cache of a mobile device carried by a user, the cache being used for storing items associated with locations in a real-world space being visited by the user; the method comprising the steps of:

- 5 (a) determining the probability of usage of at least one item in dependence on the user's progress around the space;
- (b) transforming a cached item by compression and/or degradation, to reduce the amount of cache space occupied by that item, this item being selected on the basis of the determination carried out in step (a) in respect of that item or other items.

10

63. A method of retrieving a data item to a mobile device carried by a first user visiting a real-world space, the data item being available from a service system to mobile devices of users visiting the space; the method comprising the steps of:

- 15 (a) seeking to retrieve the data item to the first user's mobile device by transfer from another mobile device in said space;
- (b) in the event that step (a) is unsuccessful, retrieving the data item to the first user's mobile device by transfer from the service system.

20 64. A method according to claim 63, wherein the data item is associated with a location in said space, step (a) being initiated as the user approaches or is at that location and including carrying out an enquiry limited to mobile devices that are, or are likely to be, near the first user or said location, to identify a mobile device, if any, holding the data item.

25 65. A method according to claim 64, wherein said enquiry is limited to mobile devices near the mobile device of the first user by having that device make the enquiry by using a short-range communications means to ask other mobile devices if they have the data item.

30 66. A method according to claim 64, wherein said enquiry is limited to mobile devices near the mobile device of the first user or near the location associated with the data item, by monitoring the locations of the mobile devices in said space.

67. A method according to claim 64, wherein said enquiry is limited to mobile devices likely to be near the mobile device of the first user by pre-defining a set of mobile devices which are associated with users belonging to the same visit group.
- 5 68. A method according to claim 63, further comprising an on-going step of keeping a record of which mobile devices, if any, hold or are likely to be holding the data item; step (a) including carrying out an enquiry limited to mobile devices that, according to said record, hold or are likely to be holding the data item.
- 10 69. A method according to claim 68, wherein said on-going step comprises tracking at least the first one of:
- transfers of the data item from the service system to a mobile device;
 - transfers of the data item between mobile devices; and
 - deletions of the data item from a mobile device.
- 15 70. A method according to claim 68, wherein said on-going step comprises at least the first one of:
- periodically making an inventory of items currently held by each mobile device;
 - recording incremental changes to the inventory of each mobile devices as items are
- 20 added / removed.
71. A method according to claim 64 or 68, wherein in step (a) said enquiry is carried out by the first user's mobile device.
- 25 72. A method according to claim 64 or 68, wherein in step (a) said enquiry is carried out by the service system at the prompting of the first user's mobile device, the service system identifying back to the first user's mobile device at least one device holding the data item where the enquiry identifies any such device.
- 30 73. A method according to claim 63, wherein multiple data items each with a respective associated location in said space are available from the service system, the method further comprising an on-going process in which said space is treated as divided into zones and,

for each zone, the service system causes the mobile devices in the zone to load data items associated with locations in that zone beyond the normal needs of the devices whereby to increase the likelihood of step (a) being successfully effected from a mobile device in the same zone as the first-user's mobile device.

5

74. A method according to claim 63, wherein multiple data items each with a respective associated location in said space are available from the service system, the method further comprising an on-going process in which said space is treated as divided into zones and, for each zone, upon a mobile device exiting the zone, it transfers the data items it holds
10 that have associated locations in the zone being exited to devices, if any, still in said zone whereby to increase the likelihood of step (a) being successfully effected from a mobile device in the same zone as the first-user's mobile device.

75. A method according to any one of claims 63 to 74, wherein a transfer effected in step
15 (a) is effected using a communications mechanism that is different to that used for a transfer effected in step (b).

76. A method of coordinated consumption of a streamed media object by first and second devices, the media object being accessible for streaming from a server, the method
20 comprising the steps of:

- (a) streaming the media object from the server to the first device and presenting it to a user of this device;
- (b) sending from the first device to the second device, during the course of step (a), data identifying the media object and a current position reached in presenting the object to
25 the user of the first device;
- (c) in response to a request from the second device, streaming the media object from the server to the second device in a separate stream to that involving the first device, and presenting the media object to the user of the second device such that normal presentation commences at a position at, or with an advance relative to, the said current
30 position indicated in step (b).

77. A method according to claim 76, wherein presentation of the media object to the user of the first device in step (a) continues whilst the second device requests and starts to receive the media object in its own stream in step (c), the second device commencing normal presentation of the media object at an advance relative to the said current position indicated in step (b) comprising at least one of the following components:

- a component taking account of the time taken to pass said data identifying the media object from the first device to the second device;
- a component taking account of the time for the second device to request and start to receive the media object stream from the server;
- a component taking account of the time to move through the media object being streamed to the second device to reach a position corresponding to the estimated then current point of consumption of the media object at the first device 31A.

78. A method according to claim 76, wherein presentation of the media object to the user of the first device in step (a) is paused whilst the second device requests and starts to receive the media object in its own stream in step (c), normal presentation of the media object being started by both devices from the said current position indicated in step (b) following the second device indicating that it is at least ready to do so.

79. A method according to claim 78, wherein in step (c) the second device indicates to the first device when it is ready to commence normal presentation of the media object from the said current position indicated in step (b) but defers doing so pending a commencement signal from the first device, the first device sending this commencement signal after having received the ready-to-commence indication from the second device and the first device resuming its own presentation of the media object in coordination with the sending of the commencement signal.

80. A method according to any one of claims 76 to 79, wherein in step (c) the server starts streaming the media object to the second device from a position at or near the current position reached by the first device in presenting the media object.

81. A method according to any one of claims 76 to 80, wherein subsequent to the commencement of presentation of the media object by both the first and second devices in at least approximate coordination, presentation coordination signals are periodically sent from at least one device to the other to enable the latter to adjust its presentation of the media object to bring it into closer coordination with the presentation by said one device.

82. A method according to any one of claims 76 to 81, wherein subsequent to the commencement of presentation of the media object by both the first and second devices in at least approximate coordination, a change signal is sent from at least one device to the other upon the first device changing its presentation position in or progression through the media object otherwise than as part of its normal progression therethrough, the said other device using this change signal to adjust its presentation of the media object correspondingly.

83. A method according to any one of claims 76 to 82, wherein in step (b) said data is sent from the first device to the second device by short-range communication means.

84. A method according to any one of claims 76 to 83, wherein in step (b) the identity of the media object and the current position reached by the first device in presenting the media object are sent separately from the first device to the second device.

85. A method according to any one of claims 76 to 84, wherein the first and second devices are mobile devices.

ABSTRACT**Methods and Arrangements Applicable to Exhibition Spaces**

5

A number of different methods and arrangements are disclosed that are suitable for use, amongst other uses, in facilitating a visit to an exhibition space or the like. In a preferred embodiment, a user 930) visiting a hall (10) is equipped with a mobile device (31) in communication with a service system (35). The mobile device (31) is arranged to deposit
10 virtual markers with the service system (35) as the user progresses around the hall (10), thereby enabling useful trail information to be built up, including by marker aggregation across multiple users. A user's mobile device (31) can also be used to provide guidance, in the form of sound cues, to guide a user around the hall (10). Furthermore, media items held by the service system (35) are associated with various locations (20-23) around the
15 hall (10) and a user (30) arriving at such a location is presented with the corresponding item or items. Preferably, the media items are pre-emptively loaded into a cache of the user's mobile device (31) in dependence on the user's progress around the hall (10). Items can also be flushed from cache on this basis though an alternative, or additional, strategy is to reduce the cache space taken up by an item by degrading it. In order to reduce load on
20 the server, a media item needed by a mobile device (31) can advantageously be first sought from a nearby device using a short-range communication mechanism. A method is also disclosed for coordinating the consumption of a streamed media object by users of, for example, the same tour party.

25 (Figure 1)

1/12





Figure 1



2/12

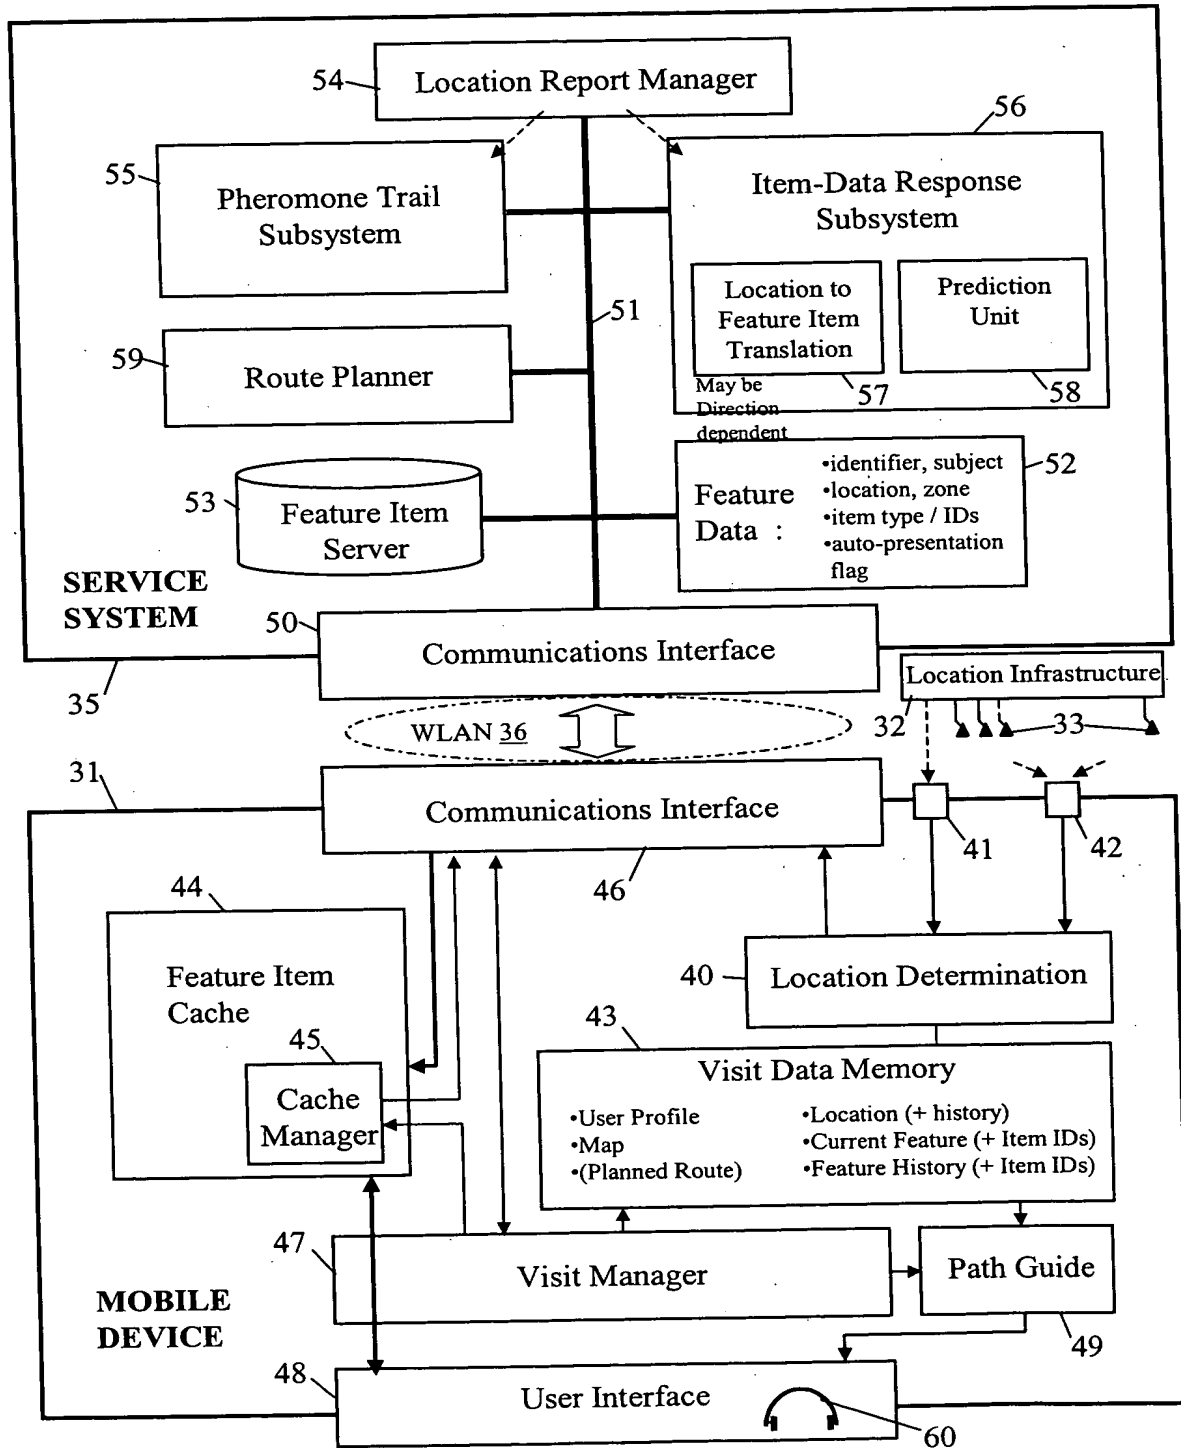


Figure 2

3/12

Location Report 62

User Identity	Location	Profile Data	Prediction-Assist Data
---------------	----------	--------------	------------------------

Figure 3

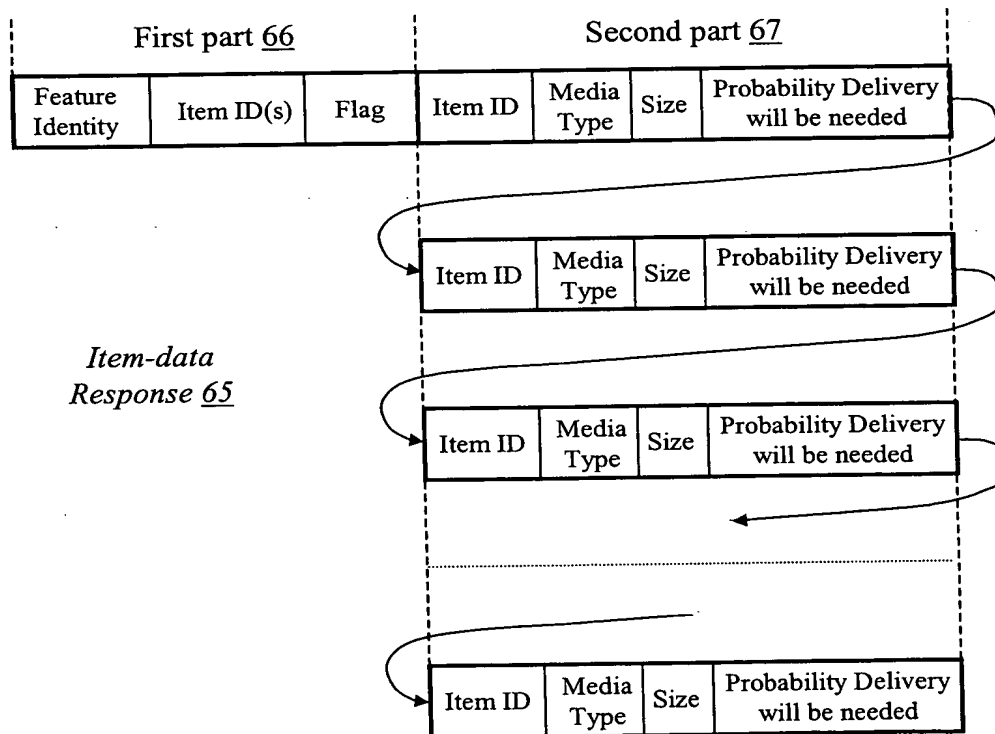


Figure 4

4/12

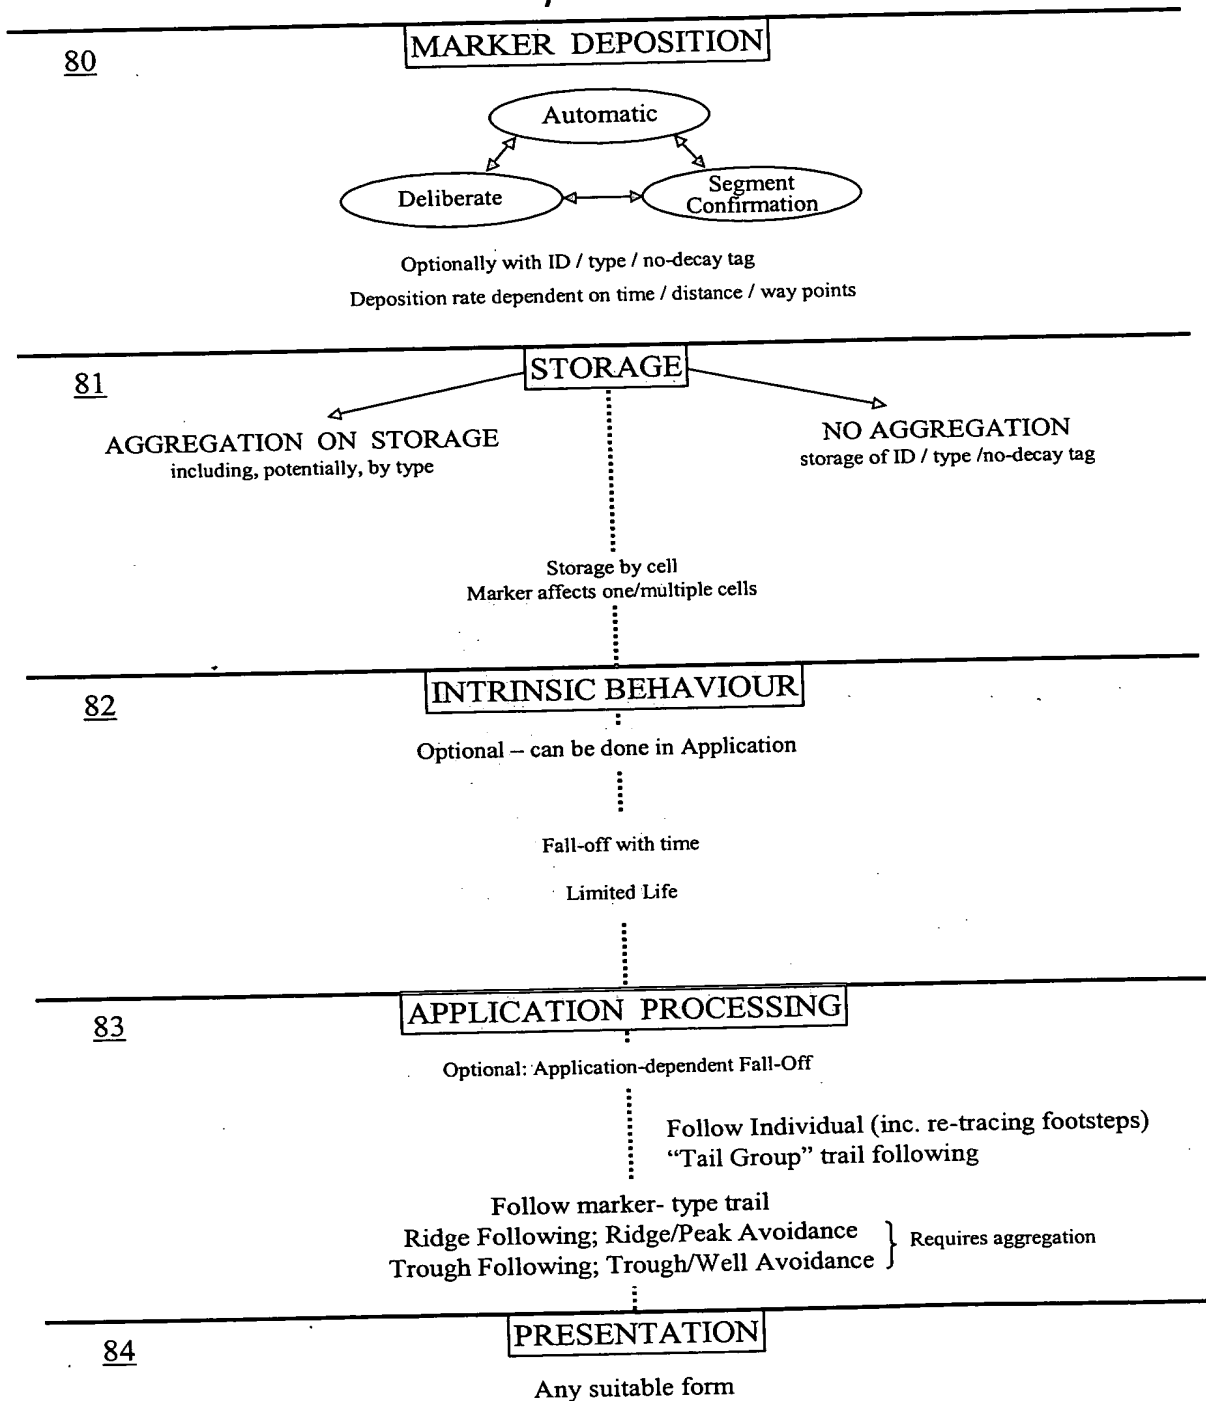



Figure 5



5/12

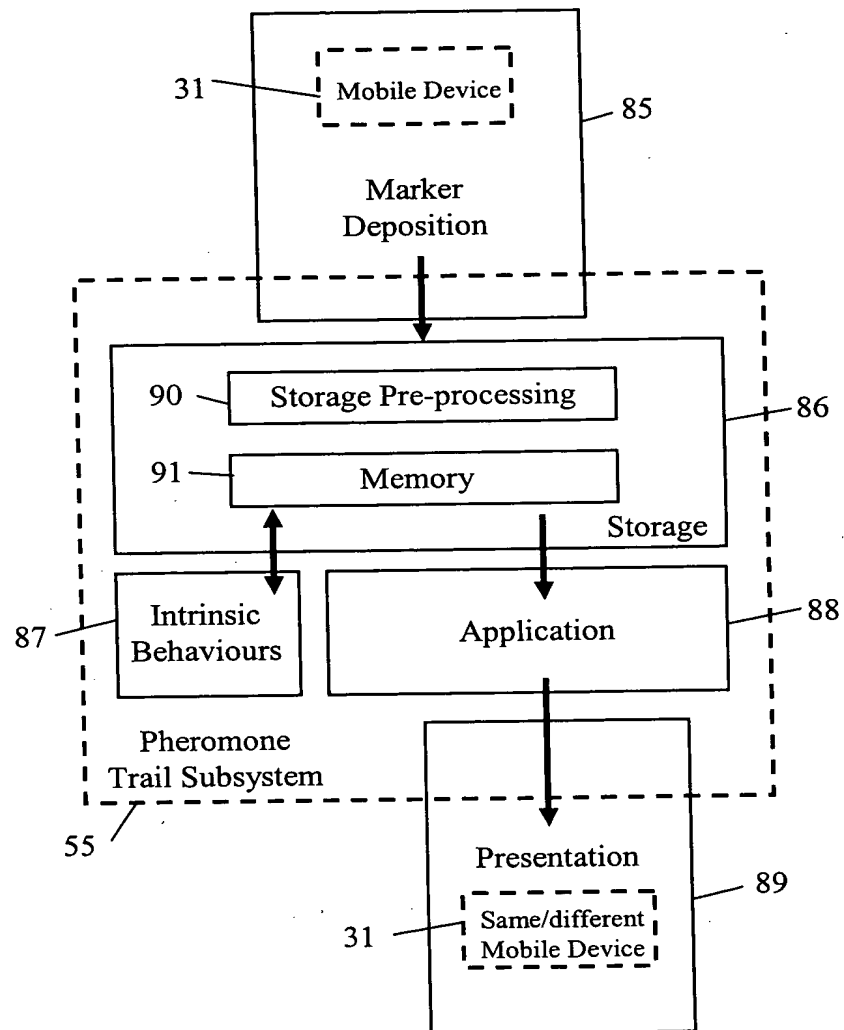


Figure 6

6/12




Figure 7




Figure 8

7/12

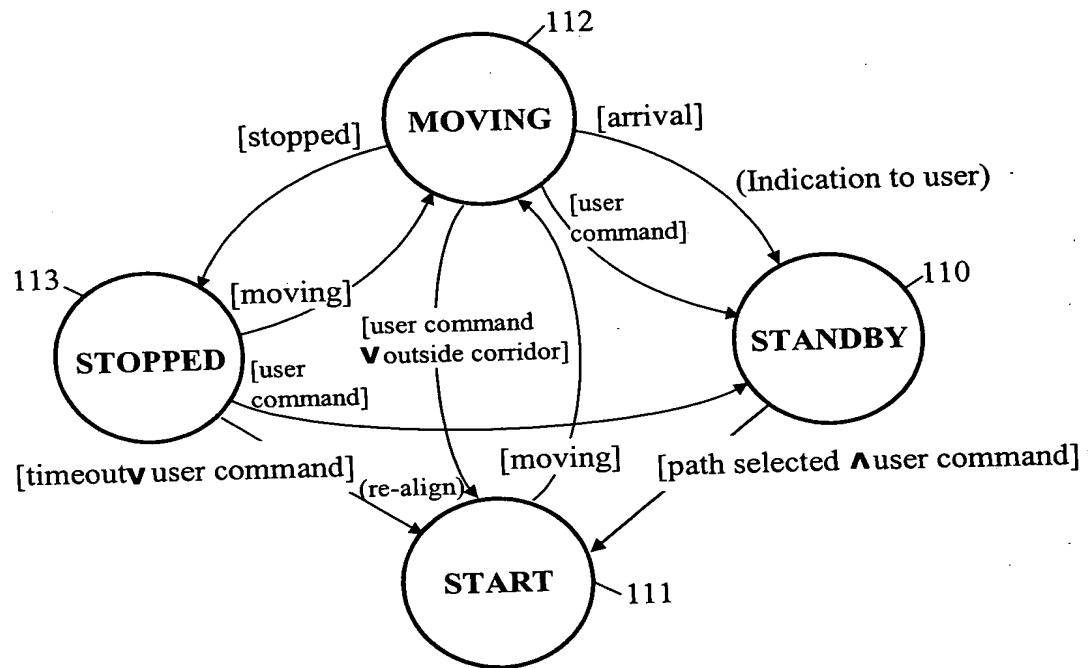


Figure 9




Figure 10

8/12




Figure 11




Figure 12

9/12




Figure 13




Figure 14

10/12

Figure 15

140

Most Recently Visited Feature

Next-Visited Feature

	F1	F2	F3	F4	F5	FN
F1	-	79	45	23	2				34
F2	42	-	56	26	10				18
F3	10	47	-	81	34				7
F4	7	9	36	-	93				0
F5	0	3	11	39	-				1
...						-			
...							-		
...								-	
FN	34	2	1	0	0				-

NEXT-FEATURE FREQUENCY MATRIX




Figure 18

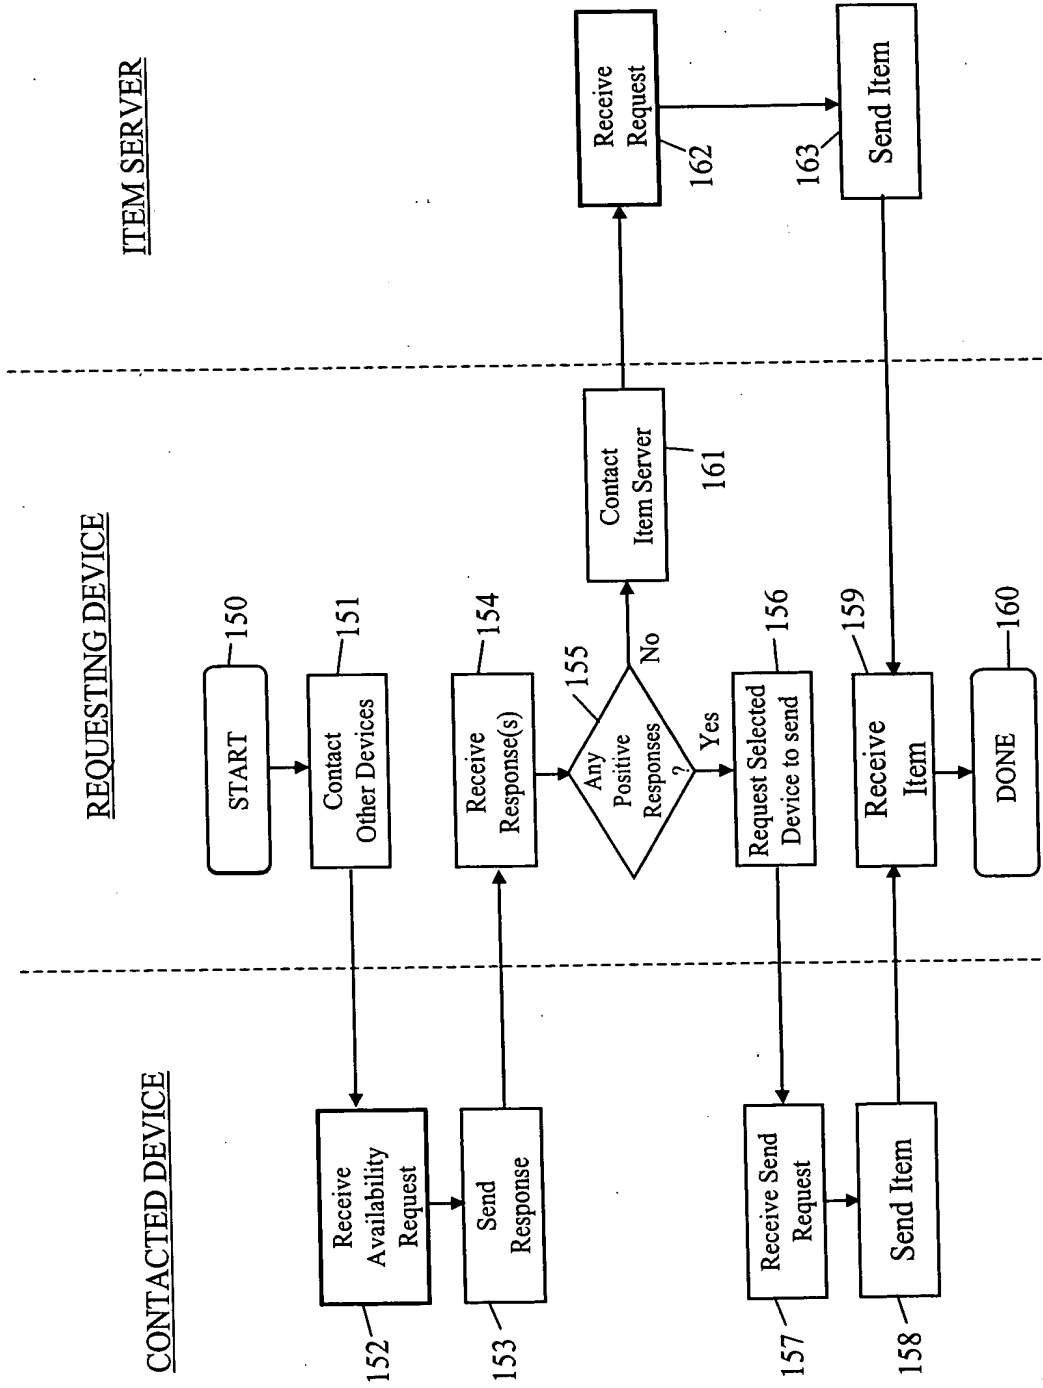


Figure 16

